

**Razor**  
Automate The Bare Metal  
Provisioning Process

**SaltStack**  
Use It To Manage  
Repetitive Tasks

₹ 100

# OpenSource

Volume: 03 | Issue: 03 | Pages: 108 | December 2014

THE COMPLETE MAGAZINE ON OPEN SOURCE

## ForYou

An **EFY** GROUP Publication

## *Cool Tools For* **WEB APPLICATIONS**

**Web App Development**  
With NoSQL Databases

■ **Enhance Java Applications**  
With FOSS APIs



11th Edition

**OPEN**  
SOURCE INDIA  
NIMHANS Convention Center

7 - 8  
November  
2014  
BENGALURU

**POST SHOW  
REPORT**



# Raise Your Expectations, Because Now There's a **Cloud** with a **Golden Lining**.

Get eNlight Cloud App Now :

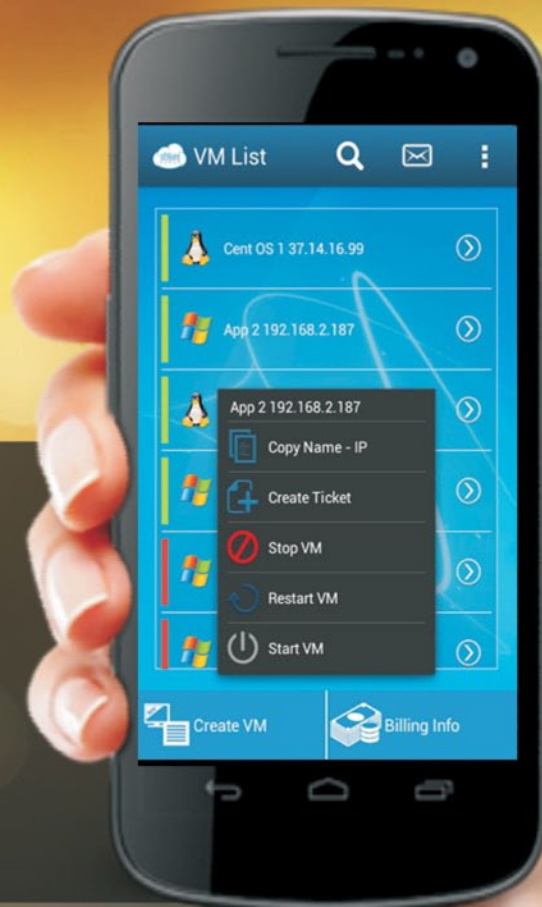


Scale, Monitor and Manage Your  
Virtual Machines anytime anywhere

Get Notification & Alert messages to  
know your Virtual Machine's Health

Get directly connected with your  
Account Manager in case of any difficulty

Activity Log to monitor every performed  
operation on your Virtual Machine





# Empower Your Business with High-End Dedicated Servers

*Deploy, Customize & Track Applications of Your Own Choice*

GET **50% OFF**

On Your First Bill

\*Valid on Quarterly, Half-Yearly & Annual Bills

Don't Miss Out On This  
**EXCLUSIVE  
OFFER!**



For more information, call us now at  
**1800-212-2022** (Toll Free) or  
**0120-666-7777** (Sales)

Fast Speed

High  
Availability

Robust  
Performance

Multi Domain  
Name  
Creation

Stringent  
Security  
Protocols

99.95%  
Network  
Uptime

Visit our website to avail  
the offers now

[www.go4hosting.in](http://www.go4hosting.in)

**Support:** +91-120-666-7718 **Email:** [sales@go4hosting.in](mailto:sales@go4hosting.in)  
**Business/Reseller Enquiry:** [sanjay.tripathi@cyfuture.com](mailto:sanjay.tripathi@cyfuture.com)

# Contents

## DEVELOPERS

- 34** An Introduction to Firefox OS
- 39** Constant Pointers and Pointers to Constant: A Subtle Difference in C Programming
- 42** Guard Your Web Applications Against the Top 10 Vulnerabilities
- 50** What is the Haskell I/O?
- 55** Java Virtual Machine (JVM) Delving Deep into its Architecture
- 58** Developing Applications Using NoSQL Databases

## ADMIN

- 61** The Need to Address Security Concerns in the Cloud Environment
- 65** Automate the Bare Metal Provisioning Process through Razor
- 69** Secure Your Linux Box With Effective User and Group Management
- 73** Automate Repetitive Tasks with SaltStack

## FOR YOU & ME

- 80** Speed Up your Research Paper with zotero
- 84** Use the History Command Effectively to Make Your Work Easier
- 88** Tweak Firefox to Get a Better Browsing Experience
- 90** Why Care About Technical Debt?



## 28 Open Source India 2014: The Industry Meets the Community



## 46 Enhance Java Applications with FOSS APIs

### REGULAR FEATURES

- |                               |                              |
|-------------------------------|------------------------------|
| <b>08</b> You Said It...      | <b>12</b> FOSSBytes          |
| <b>09</b> Offers of the Month | <b>72</b> Editorial Calendar |
| <b>10</b> New Products        | <b>104</b> Tips & Tricks     |



# Open Source For You

## Android application

is now available for FREE on



Download it now



For any further queries

Connect to us on



OpenSourceForU

OR write to

omar.farooq@efyindia.com



# Contents



## 76 Playing Around with Graphs in Maxima



- 98** "IT requirements, business constraints and workload requirements determine the best choice for a server"—  
*Gopalswamy Srinivasan, category manager, servers, HP India*



- 100** "Mozilla Smartphones are Here to Replace Feature Phones!"—  
*Dr James Ho, senior director of mobile devices, Mozilla*

## OPEN GURUS

- 93** Develop a GNU/Linux-like OS for a Single Board Computer
- 96** Install BURG, the Brand-new Universal loader from GRUB
- 17** Indian Government Opts for Open Source in New Policy!
- 18** Microsoft Releases Source Code of .NET for Linux and Mac OS!

## COLUMNS

- 24** CodeSport
- 26** Exploring Software: Exploring Big Data on a Desktop, with Hadoop, Elasticsearch and Pig

## Buyers' Guide

How to Select the Right Desktop Scanners **20**

## DVD Of The Month

Try your hands at some of the great Linux distros available.

- Ubuntu 14.10 • Ubuntu MATE 14.04
  - Ubuntu Server 14.10 • openSUSE 13.2 KDE Live
  - openSUSE 13.2 GNOME Live
- 106**

## Editor

RAHUL CHOPRA

## Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020

Ptc: (011) 26810602, 26810603; Fax: 26817563

E-mail: info@efy.in

## Missing Issues

E-mail: support@efy.in

## BENGALURU

Ptc: (080) 25260394, 25260023

E-mail: efbtl@efy.in

## Customer Care

E-mail: support@efy.in

## Back Issues

Kits 'n' Spares

New Delhi 110020

Ptc: (011) 26371661, 26371662

E-mail: info@kitsnspares.com

## Advertising

### CHENNAI

Ptc: (044) 42994363

E-mail: efyeng@efy.in

### HYDERABAD

Ptc: (040) 67172633

E-mail: efyeng@efy.in

### KOLKATA

Ptc: (033) 22294788

E-mail: efyeng@efy.in

### MUMBAI

Ptc: (022) 24950047, 24928520

E-mail: efmum@efy.in

### PUNE

Ptc: (020) 40147882

E-mail: efpune@efy.in

### GUJARAT

Ptc: (079) 61344948

E-mail: efyahd@efy.in

### JAPAN

Tandem Inc., Ptc: 81-3-3541-4166

E-mail: tandem@efy.in

### SINGAPORE

Publicitas Singapore Pte Ltd

Ptc: +65-6836 2272

E-mail: publicitas@efy.in

### UNITED STATES

E & Tech Media

Ptc: +1 860 536 6677

E-mail: veronique@amarque@gmail.com

### CHINA

Power Pioneer Group Inc.

Ptc: (86 755) 83729797, (86) 13923802595

E-mail: powerpioneer@efy.in

### TAIWAN

J.K. Media, Ptc: 886-2-87726780 ext. 10

E-mail: jkmedia@efy.in

## Exclusive News-stand Distributor (India)

IBH BOOKS AND MAGAZINES DISTRIBUTORS LTD

Unit No.10, Bezzola Complex, Opp. Suman Nagar,

Sion Trombay Road, Chembur,

Mumbai - 400 071

Phones: 022 - 40497401/02

E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Tasa Art Printers Pvt Ltd, A-46/47, Sec-5, Noida, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2013. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-NonCommercial 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-nc/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

## SUBSCRIPTION RATES

| You Pay |      | Overseas |
|---------|------|----------|
| Year    | (₹)  | (₹)      |
| Five    | 6000 | 3600     |
| Three   | 3600 | 2520     |
| One     | 1200 | 960      |
|         |      | US\$ 120 |

Kindly add ₹ 50/- for outside Delhi despatch.

Please send payments only in favour of EFY Enterprises Pvt Ltd. Non-receipt of copies may be reported to support@efy.in—do mention your subscription number.





Address Book



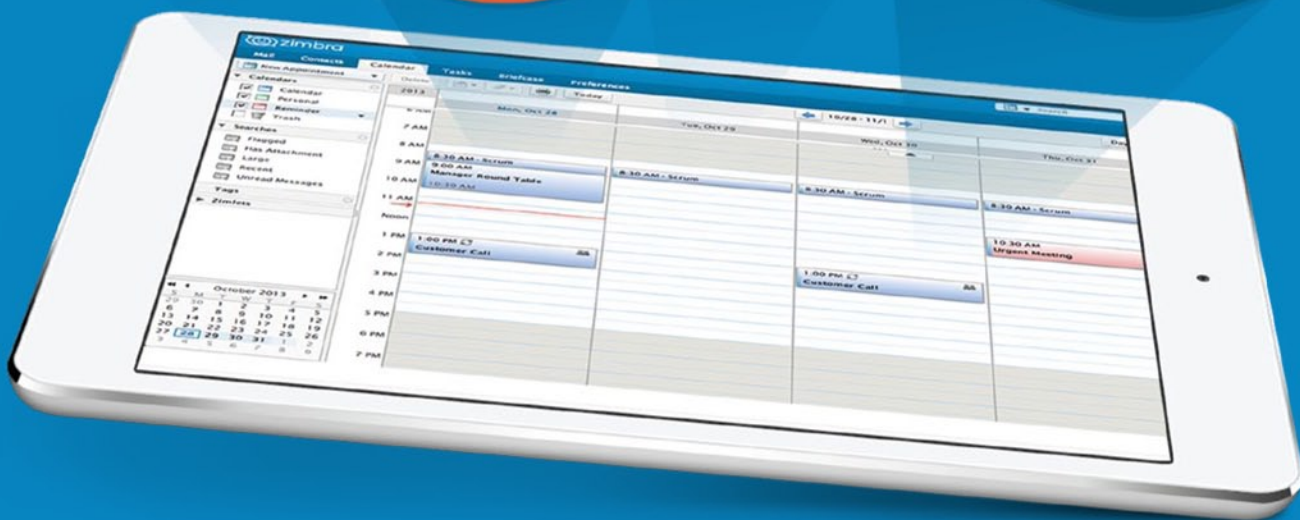
Email



Calendar



Communications



# Zimbra Collaboration



**Zimbra Collaboration offers all of the above and more!**

- ✓ Active Sync support enabling complete mobile device management
- ✓ Native backup & restore tools
- ✓ File sharing with access & version control
- ✓ Archiving & discovery tool for compliance & audit purposes



Contact sales at **+91 20 6730 4731**

or email **[india.marketing@zimbra.com](mailto:india.marketing@zimbra.com)** for further information

Free Trial, visit  
[www.zimbra.com](http://www.zimbra.com)





# YOU SAID IT



## A thumbs-up for the OSFY DVD

I bought the November issue of OSFY and I loved the FreeBSD OS. I just want to thank the OSFY team for providing such useful operating systems to its readers so that we may try out the latest technologies open source offers.

—**Parveen Kumar**  
[parveen199214@gmail.com](mailto:parveen199214@gmail.com)

**ED:** Thanks for the compliment, Parveen. It's good to know that you've been a regular reader and like our DVDs so much.



## Buying previous issues of OSFY

I would love to get previous issues of OSFY (from January 2013 to October 2014). Is there a way I can buy these from you? I would really appreciate it if I could!

—**Sudheer Vaidya**  
[sudheervaidya@yahoo.co.in](mailto:sudheervaidya@yahoo.co.in)

**ED:** Sure, you can definitely get all the previous issues of OSFY! All you need to do is to write to our support team at [support@efy.in](mailto:support@efy.in)



## In search of code review tools

I am looking for open source code review software that can be installed on Windows servers for free with an integrating SVN. It would be a boon if it integrates with Visual Studio.

—**Mallikarjun S Bendigeri**  
[mbendigeri@gmail.com](mailto:mbendigeri@gmail.com)

**ED:** Thank you for your query. We have a Facebook page called Open Source For You for all the open source enthusiasts. You can post your query on the page and we are sure the community will respond with the right solutions.



## From an ardent follower of OSFY's Facebook page

I am a big follower of the Facebook page 'Open Source For You' as it offers a lot on open source. I am new to computer engineering and I want to know more about the Linux operating system. Can

you suggest how I can learn to excel in programming while using open source. That would be a great help.

—**Mearg Taddese**

**ED:** Thank you for 'liking' our Facebook page. We wish to continue sharing important content related to FOSS. You may post this query on our page and you will definitely get a suitable reply from the open source community.



## Suggestion for the 'Tips & Tricks' section

I am an avid reader of OSFY and I love the contents of the magazine, particularly the ideas you share in the section 'Tips & Tricks'. These are quite useful and interesting and I never fail to implement them. Keeping this in mind, I have a suggestion: why don't you introduce a new section based only on Linux recipes. It will be a success.

—**Ashish Kumar**  
[ashish.kumar11@outlook.com](mailto:ashish.kumar11@outlook.com)


**ED:** Thank you, Ashish, for such valuable feedback. It is always a treat to know that people are not just reading our magazine but gaining knowledge from it. We will surely give a thought to introducing a section on Linux recipes. Till then, keep reading OSFY!



## On writing for OSFY

I love reading OSFY and now it has inspired me to also contribute content to your magazine since I, too, am an open source enthusiast. Can you please explain the procedure and also the eligibility criteria for contributing content to your magazine?

—**Neha Vashishth**  
[neha\\_v9@yahoo.com](mailto:neha_v9@yahoo.com)

**ED:** It's great to hear of your interest in writing for us! It feels good when enthusiastic readers like you wish to pen an article for us. Before doing so, you can send us a detailed 'Table of Contents' (ToC) at [osfyedit@efy.in](mailto:osfyedit@efy.in) for the topic you want to write on. Our team will review it and once they give you the thumbs-up, you can go ahead with the article. Please feel free to get in touch if you need anything else from us.  **END**



## Please send your comments or suggestions to:

The Editor,  
Open Source For You,  
D-87/1, Okhla Industrial Area, Phase I,  
New Delhi 110020, **Phone:** 011-26810601/02/03,  
**Fax:** 011-26817563, **Email:** [osfyedit@efy.in](mailto:osfyedit@efy.in)



# OFFERS

# THE MONTH

## One month free



**Free Dedicated Server Hosting for one month**

Subscribe for our Annual Package of Dedicated Server Hosting & enjoy one month free service

**Hurry!**  
Offer valid till 31<sup>st</sup> December 2014!

For more information, call us on 1800-209-3006/ +91-253-6636500

[www.esds.co.in](http://www.esds.co.in)

## 2000 Rupees Coupon

(Free Trial Coupon)



**No condition attached for trial of our cloud platform**

Enjoy & Please share Feedback at [sales@cloudoye.com](mailto:sales@cloudoye.com)

**Hurry!**  
Offer valid till 31<sup>st</sup> December 2014!

For more information, call us on 1800-212-2022 / +91-120-666-7718

[www.cloudoye.com](http://www.cloudoye.com)

## Get 10% discount



**Reseller package special offer !**

Free Dedicated hosting/VPS for one month. Subscribe for annual package of Dedicated hosting/VPS and get one month FREE

**Hurry!**  
Offer valid till 31<sup>st</sup> December 2014!

Contact us at 09841073179 or Write to [sales@space2host.com](mailto:sales@space2host.com)

[www.space2host.com](http://www.space2host.com)

## 35% off & more



**"Do not wait! Be a part of the winning team"**

Get 35% off on course fees and if you appear for two Red Hat exams, the second shot is free

**Hurry!**  
Offer valid till 31<sup>st</sup> December 2014!

Contact us @ 98409 82184/85 or Write to [enquiry@vectratech.in](mailto:enquiry@vectratech.in)

[www.vectratech.in](http://www.vectratech.in)

## Get 12 Months Free



**Pay Annually & get 12 Month Free Services on Dedicated Server Hosting**

Subscribe for the Annual Packages of Dedicated Server Hosting & Enjoy Next 12 Months Free Services

**Hurry!**  
Offer valid till 31<sup>st</sup> December 2014!

For more information, call us on 1800-212-2022 / +91-120-666-7777

[www.goforhosting.com](http://www.goforhosting.com)

## Get 25% Off



**Considering VPS or a Dedicated Server? Save Big !!! And go with our ProX Plans**

25% Off on ProX Plans - Ideal for running High Traffic or E-Commerce Website  
Coupon Code : OSFY2014

**Hurry!**  
Offer valid till 31<sup>st</sup> December 2014!

Contact us at 98769-44977 or Write to [support@packwebhosting.com](mailto:support@packwebhosting.com)

[www.prox.packwebhosting.com](http://www.prox.packwebhosting.com)

## Fees Rs. 15000/- (SC/ST/PH-50% Concession)



**BIG DATA ANALYTICS**  
using hadoop eco-system

Venue: Malleshwaram, Bangalore  
Qualification: Degree / Diploma  
Training Contents: Introduction, HDFS, Mapreduce, Administration, Hbase, Hive, Pig, Sqoop, Flume, Oozie, Hadoop2.0

**Certificates Will Be Awarded**

For more information, call us on Balajee. K.M, Asst, Director, Mob.: 9900226980, 080 41100557

[www.msmedibangalore.gov.in](http://www.msmedibangalore.gov.in)

## Register in Advance to receive a special discount

**EMBEDDED SOFTWARE DEVELOPMENT COURSES AND WORKSHOPS**

**School of Embedded Software Development**

Faculty - Babu Krishnamurthy  
Visiting Faculty/ CDAC/ACTS Porting  
Embedded RTOS to an ARM platform

Dates : November 2014 & January 2015  
Date will be on Request

CDAC Alumni will get a special rate if registered in advance -please write to [babu\\_krishnamurthy@yahoo.com](mailto:babu_krishnamurthy@yahoo.com) mentioning the CDAC/ACTS center and batch details



# NEW PRODUCTS

## Here's a Bluetooth-enabled keyboard to switch between paired devices



Logitech has introduced an innovative Bluetooth keyboard, the K480, which is compatible with multiple devices. It's a wireless gadget that works with Mac, Windows and Android devices. The selection dial, located above the Escape key, allows switching between paired devices.

The key features of the K480 are that it supports up to three different Bluetooth devices, enables easy switching between multiple device profiles and also saves battery life. It is cross-platform and there is no need for a dongle, as long as the target device is Bluetooth-enabled. The function keys have some multimedia or mobile-specific buttons and for Android devices, an extra app can be installed to provide additional keyboard layouts. The keyboard has two pre-installed AAA batteries. The best feature of the keyboard is definitely its key layout, which is quite spacious and offers a quick set-up process at a modest price, for almost all major operating systems.

**Price: ₹ 3090**

**Address:** Logitech Electronics India Pvt Ltd, 601, Raja House, 30-31, Nehru Place, New Delhi; **Ph:** 011-47306602; **Website:** <http://www.logitech.com>

## The Oppo R5 could be the world's slimmest smartphone

Oppo Electronics, the Chinese firm, has launched the Oppo R5 handset, which is touted to be the world's slimmest smartphone. The Oppo R5 smartphone is just 0.48 cm thick and has surpassed the Gionee Elife S5.1 handset, which is 0.5 cm thick. It was even recognised by the Guinness Book of World Records as the slimmest smartphone when it was launched. It has got both metallic and ceramic finishes at the corners, and supports a 13.21-cm full HD AMOLED display with a 1080 x 1920 pixel resolution and 423 ppi pixel density. It is powered by a 64-bit octa-core Qualcomm Snapdragon 615 (MSM8939) SoC clocked at quad 2.1GHz and quad 1.5GHz. It has 2GB RAM, Adreno 405 and 16GB internal storage memory, which are not further expandable. The device runs on the company's proprietary ColorOS 2.0.1 (based on Android 4.4 KitKat) and supports a single micro-SIM. The Oppo R5 smartphone sports a 13MP rear camera and a 5MP front-facing shooter. The primary camera features a Sony IMX214 sensor, an f/2.0 aperture and an LED flash. It has a 2000mAh battery with VOOC technology, which can charge the device from 0 per cent to 75 per cent in just 30 minutes.



**Price: ₹ 30,500**

**Address:** Oppo, Vatika Business Park, 2nd Floor, Sohna Road, Block 1, Gurgaon; **Ph:** 18001032777; **Website:** <http://oppomobile.in/>

## Intex releases the Aqua Amaze, with Android KitKat

Intex has released its Aqua Amaze smartphone in India in two colour variants, light blue and black. The Aqua Amaze comes with dual-SIM (WCDMA+GSM) support and runs on Android 4.4.2 KitKat. The device features a 12.7-cm HD IPS display with a 540 x 960 pixel resolution and is covered with OGS lamination. It is powered by an octa-core (MediaTek MT6592M) processor clocked at 1.4GHz, and packs in 1GB RAM and 8GB internal storage space, which can further be expanded up to 32GB via microSD card. The device supports a 13MP rear autofocus camera with a flash and a 5MP front shooter. The rear camera features many add-on options like panorama, face beauty, smile shot, face detection, HDR, self-timer, tap-to-capture, and voice capture.




**Price: ₹ 10,690**

**Address:** Intex Technologies (India) Limited, D - 18/2, Okhla Industrial Area, Phase - II, New Delhi -110020; **Ph:** 91 11 41610224/25/26, **Email:** [info@intextechnologies.com](mailto:info@intextechnologies.com); **Website:** <http://www.intextechnologies.com>



## Dell launches two Android tablets in India

Dell has introduced the refreshed versions of its Venue 7 and Venue 8 Android tablets in India. Both run the Android 4.4 KitKat OS and are available with 16GB internal storage space, which can further be expanded up to 64GB via microSD card. Also, both the devices are backed by a 4550mAh battery. The Venue 7 tablet sports a 17.78-cm IPS display WXGA with a 1280 x 800 pixel resolution, and is powered by the 1.6GHz dual-core Intel Atom Z3460 'Merrifield' processor along with 1GB RAM. The device has a 5MP rear camera and a 1MP front-facing camera. The Venue 8 tablet has a 20.32-cm (8-inch) full-HD IPS display with a 1920 x 1200 pixel resolution. It is powered by a 2.1GHz dual-core Intel Atom Z3480 'Merrifield' processor coupled with 1GB RAM. It comes with a 5MP rear camera and a 2MP front-facing camera.



**Price of the Venue 7 : ₹ 11,999**

**Price of the Venue 8 : ₹ 14,999**

**Address:** Dell India Pvt Ltd, Nos 12/1, 12/2A and 13/1A, Divya Shree Greens, Koramangala Ring Road, Bengaluru, Karnataka 560036; **Ph:** 08025357311, **Email:** info@dell.in; **Website:** www.dell.com

## Raspberry Pi Foundation launches cheapest ever miniature computer—the Model A+

A new and smaller version of Raspberry Pi Model A has just been unveiled by the Raspberry Pi Foundation, and it has been named Model A+. Its basic configuration is similar to the original Model A. Some features have been adopted from the recently launched Model B+. Model A+ is packed with a Broadcom BCM2835 SoC and 256MB RAM, just like its predecessor, but the SD card slot has been switched to a microSD slot. There are 14 new general-purpose input-output (GPIO) pins too, which can be used for building connections between additional sensors. The audio circuit has also been improved, which ensures a low-noise power supply. The board also consumes less power, and its width has been reduced to 65 mm from 86 mm.



The features of the Broadcom BCM2835 SoC are its 700MHz ARM11 CPU core and Vidcore IV GPU. Its video output is supported by HDMI. This is the cheapest ever miniature computer from the Raspberry Pi Foundation, making it more affordable for students and budding entrepreneurs to build exciting devices. It was first introduced along with Model B+. The Raspberry Pi Model A+ boards are getting manufactured at a Sony-owned factory in the UK.

**Price: ₹ 1,230**

**Address:** Kits'n'Spares, D 88/3, Basement, Okhla Industrial Area, Phase-1, New Delhi-110020; **Ph:** 01126371661; **Website:** http://kitsnspares.com; **Email:** info@kitsnspares.com

## Zebronics introduces two new external power banks



Zebronics, has introduced external power banks in two capacities to keep smartphones and tablets running all the time. The Zeb-PG4400, with a capacity of 4400mAh, and the Zeb-PG2200, with a capacity of 2200mAh, are ideal for today's generation of 'connected' mobile enthusiasts, whether they are working, talking, surfing or sharing content on social media.

Available in two elegant colours, smoked grey and lively orange, the device has a soft-touch, scratch-resistant rubber finish and is ultra-light —weighing just 120 grams. In use, it provides a standard USB port for charging most compatible smartphones and tablets via the included matching micro-USB cable, for uninterrupted mobility. Its in-built LED torch can be handy in an emergency. An LED indicator on the Zeb-PG4400 external power bank shows the device's readiness to charge or the amount of charge left in the charger. Safety features include input and output short protection, over-charging and over-discharging protection and an auto-recovery feature after short protection.

**Price of the Zeb-PG4400: ₹ 900**

**Price of the Zeb-PG2200: ₹ 400**

**Address:** Zebronics, Eros Building - 56, 610, Nehru Place, New Delhi, Delhi 110019; **Ph:** 011 3192 0333; **Website:** www.zebronics.com/

## Indonesia tax agency saves on expenses by using open source software!

According to a latest report published by the Indonesian Tax Agency, it has saved over 90 per cent on expenses because of the adoption of open source systems.

Low maintenance and the free licences of open source software play a major role in such high savings. The agency adopted open source software for most of the applications in 2002. Only some critical systems that need constant support run on proprietary software.

The complexity of the procurement process in Indonesia is another reason why the agency has adopted open source software, as open source doesn't come under a procurement process in the country. The agency is still using proprietary software for document management, but is planning to migrate to an open source document management system very soon.



## An open source foldable 3D printer has arrived!

3D printers help innovators and manufacturers to quickly build prototypes of their products. The industry is adopting 3D printing technology pretty rapidly. So, it's no wonder that innovators are coming up with better models of 3D printers. The latest TeeBotMax 3D printer is an example of such innovation—

it is the world's first foldable 3D printer. Open source developers have built it. The printer is very simple to build, easy to transport and does not compromise on printing quality. TeeBotMax comes with a DIY guide to build the printer, for which a user doesn't need to be an expert.



The detailed PDF guide is available on open source community sites. There are even how-to YouTube videos to help you out. The material required to build TeeBotMax can be ordered from [www.3dstuffs.nl](http://www.3dstuffs.nl) or you can even pick it from the local hardware store. The printer is made up of an aluminium square pipe frame. The PDF guide includes a list of the required supplies. The printer is just 22.86-cm (9-inch) tall when folded and can produce up to a 14 cm high vase, for which it requires two-and-a-half hours. Flexible filaments can be used as material to print objects. However, this little printer makes a lot of noise while printing.

## UC Browser 10.0 available for download for Android

UC Browser 10.0 is now available for Android, and it can be downloaded via Google Play. This latest version offers a new UI, which has been adopted from Google's Material Design, which recently went live with Android 5.0 Lollipop. UC Browser 10.0 for Android offers simple settings and the navigation page has also undergone

## Latest version of Red Hat Software Collections now available

software  
COLLECTIONS

Red Hat has announced the general availability of Red Hat Software Collections 1.2, which delivers the latest, stable versions of essential development tools, dynamic languages and open source databases, all on a separate lifecycle from Red Hat Enterprise Linux. The third instalment of Red Hat Software Collections now includes vital open developer tools such as GCC 4.9 and Git, and, for the first time, makes the Eclipse IDE available on Red Hat Enterprise Linux 7. In addition, Red Hat is offering Dockerfiles for many of the most popular software collections, aiding in the rapid creation and deployment of container-based applications. New additions to Red Hat Software Collections include: Red Hat Developer Toolset 3 as a software collection, bridging development agility and production stability with the latest stable versions of essential open development tools, including ltrace 0.7.91 and a number of other enhancements; devAssistant 0.9.1—a tool for setting up development environments and publishing code; Maven 3.0—a build automation tool for Java projects that describes how software is built and the associated dependencies; Git 1.9.4, which is now separate from the Red Hat Developer Toolset in its own collection that allows developers to access and install Git as a standalone, and many more.



redesigning. The new version of UC Browser follows version 9.5 which offered faster speeds, a Web application centre and an image viewer. Now, this latest version comes along with several bug fixes and improved management of tabs. The application also comprises a new dynamic gesture control feature, and it allows users to switch between tabs quite easily. Other highlighted features include incognito browsing, smart downloading, custom themes, and so on. The updated browser is a 13MB download and it needs Android 2.2 or higher to work.

## Canonical releases Ubuntu 14.10

Canonical has just brought out a major update to its desktop OS. Ubuntu 14.10 comes with a new and improved desktop for the desktop version as well as



important improvements in its cloud and server editions. The company has introduced secure hypervisors and container technology along with Docker v1.2 in the server and cloud editions of Ubuntu 14.10. Canonical claims that the user can control the container without any super-user

authentication in this latest update. Users will be able to run more applications on the same server or cloud using container support. This will replace currently used technologies such as KVM. Ubuntu 14.10 also supports many cloud technologies such as Hadoop, Hive, Elasticsearch, PigLatin and Apache Storm. Ubuntu is highly integrated to support Infrastructure as a Service (IaaS) to enable small and large scale businesses to scale up and down their data centre and server capacities, as required. Canonical claims that this is in the best interests of all connected businesses. This new release brings the high-level DevOps tool, Juju, to Ubuntu. It enables developers to easily deploy and scale applications on the cloud or on bare metal, and helps them to scale out workloads at the machine level and the service level.

## Zentyal Server 4.0, a major new Linux release for small business servers

Zentyal, the developer of server technology natively interoperable with Microsoft server products, has announced a new release of the Zentyal Linux small business server. Zentyal Server 4.0 aims at offering small and medium businesses (SMBs) a Linux-based small business server that can be set up in less than 30 minutes and is both easy-to-use and affordable. The Zentyal 4.0 release focuses on providing an improved and stable server edition with native Microsoft Exchange protocol implementation and Active Directory interoperability. The aim is to provide easy-to-use small business servers, with native support for mixed IT environments that include Windows, Linux and Mac OS clients, as well as mobile devices with ActiveSync. Besides focusing primarily on mail and mail-related directory features, additional improvements have also taken place. The L2TP module has been restructured and improved, and free configuration backup in the cloud has been made available directly through the Zentyal Server UI. Moreover, major efforts have been put into establishing the necessary quality assurance processes, to improve the stability of the Zentyal Server releases.



## Seagate supports new version of Microsoft's Open CloudServer

Seagate Technology has announced support for Microsoft's Open CloudServer version 2. This version will include performance enhancements and expand the management software code provided to



the open source community, including new deployment and diagnostics functions. The enhanced version will also simplify deployment while enabling greater flexibility and lowering implementation costs. Open CloudServer version 2 will allow for improved storage solutions while enabling changes to match the dynamic cloud environment. The specification will help optimise storage solutions for large, Web-scale deployments by allowing for greater flexibility while reducing complexity. The costs of storage solutions can also be reduced through the elimination of cabling as the power management is distributed through the backplane. These improvements are part of an ongoing movement towards optimisation in cost, performance and the implementation of cloud infrastructures. The Open CloudServer specification is part of continuing work that Seagate and Microsoft participate in to share cloud technologies and experiences with the Open Compute Project industry group. The group collaborates to define and promote open source standards for cloud computing. The goal is to help cloud builders develop more customisable solutions by using open platforms, while reducing operating costs and providing benefits for consumers in the marketplace.

## Mozilla releases Firefox Developer Edition

Marking the 10th anniversary of the Firefox browser, Mozilla has released a new developer edition of Firefox, for Web developers, which will enable them to boost their productivity across multiple devices and platforms. The Firefox Developer Edition has a familiar UI and features, and the biggest change in it is its dark UI, which has been applied to the entire browser. If any developer is uncomfortable with this change, he can revert back to the familiar light theme through the *Customise* options, which are available at the bottom of the 'hamburger' menu. There, developers need to click on 'Use Firefox Developer Edition Theme'.



The key features of this first Developer Edition are WebIDE and Valence. The former is aimed at offering an integrated editor so that developers can work with Firefox OS apps in simulators, Firefox OS devices and also Firefox-for-Android devices. There is a debugger in the edition too, which allows one to pause and inspect components. Valence, on the other hand, supports WebIDE and allows connecting with Chrome-for-Android devices, Safari and iOS. It also lets developers inspect content on Chrome-for-Android devices, like debuggable WebView content in the KitKat OS and above. The new Firefox Developer Edition can be got from the Mozilla website itself; this edition is based on Firefox 35.0a2.

## IBM expands cloud centre network


IBM has announced the opening of a 2,790 sq m cloud centre in Airoli, Mumbai. This new facility is the latest IBM cloud centre in one of the world's major markets and is part of the company's global expansion drive in this domain. The centre further expands IBM's global cloud footprint, which includes centres in London, Amsterdam, Beijing, Hong Kong, Singapore, Melbourne, Toronto, Dallas, Raleigh and Paris. Adoption of cloud computing worldwide is driven by the need for in-country data location to meet requirements for data sovereignty—a new form of customer demand. About 100 nations and territories have already adopted laws that dictate how governments and private enterprises handle personal data. As more data moves to the cloud, nations and territories worldwide are adopting data management laws to protect data held by businesses and governments.

## Apple iPad has a LibreOffice-based document editor

LibreOffice has extended its reach now. A US-based company, CloudOn, has reportedly introduced a document editor for Apple iPad that is based on LibreOffice, which is free and open source. An official statement from the company has announced that the new app offers excellent features for creating and editing documents with a gesture-first document editor, which allows users to create new designs. The same company integrated Microsoft Office to mobile devices for the first time, which expanded mobile productivity for millions of people, said CEO and co-founder of CloudOn, Milind Gadekar. The gesture-first design allows more speed while creating and editing any document, and the app is compatible with MS Word, Dropbox, OneDrive, Google Drive and other cloud services.

## Red Hat collaborates with SAP for next-gen mobile database development

Red Hat has announced new results of its collaboration with SAP to drive the next generation of mobile database development and synchronisation. Developers now have access to new components of the SAP data management



portfolio in OpenShift, Red Hat's award-winning Platform-as-a-Service (PaaS) offering. The mobile data management cartridge for the SAP SQL Anywhere suite makes it easier for developers creating MySQL-based database applications to extend their data to remote and mobile applications, and keep it synchronised. Red Hat and SAP are working together to provide a unique universal framework of database and application development technologies geared towards supporting a wide range of operating systems, languages and services, all within a cartridge, so users can quickly gain a competitive advantage and a return on their technology investments. With the cartridge for SAP SQL Anywhere on OpenShift, developers can rapidly develop database-powered, high-performing applications through automated workflows and tooling to help them access what they need, when they need it. The combination of OpenShift by Red Hat and SAP SQL Anywhere gives developers a choice-driven solution that is capable of handling large amounts of data. OpenShift



automates the provisioning management and scaling of the application, while the self-administering, self-tuning, self-healing and synchronisation support features of SAP SQL Anywhere make it well-suited for zero-administration environments.

## Ubuntu MATE 14.04.1 LTS released

Ubuntu has released the 14.04.1 LTS version of its MATE desktop distribution. Ubuntu has been providing support for the LTS versions of its distributions for many years. The Ubuntu developers' team follows a fixed release cycle. According to the cycle, 14.04 LTS should have been released in April 2014, but the developers missed the deadline. So, Ubuntu released MATE 14.10 as per schedule in October



this year. The version number might confuse users and look older to them but MATE 14.04.1 LTS is much improved compared to MATE 14.10. Ubuntu developers have integrated many new features in 14.04.1 LTS. This new distro comes with indicators support, new panel layout, improved

accessibility features and the latest Orca 3.14. The UI of MATE 14.04.1 LTS is much more polished than 14.10. Lots of new Debian packages have been added. A lot of users have been seen adopting the Ubuntu MATE distro this year. With the new LTS version, Ubuntu users can stick to the distro for a long time and not worry about updates. Ubuntu has pushed most of these changes to MATE 14.10 via updates. Ubuntu developers will support Ubuntu MATE 14.04 LTS till 2019. Users need not worry about security fixes and important updates till 2019.

## Halcyon Software introduces a monitoring solution for Linux

Halcyon Software has announced the immediate availability of Network Server Suite's new component, the Linux on Power agent, which runs on IBM Power Systems. This new agent will ensure that companies have tighter control over key business processes and 'mission-critical' applications running on Linux, through continuous monitoring and automatic management. Halcyon's new monitoring technology meets the requirements of organisations deploying Linux on IBM Power Systems to give greater scalability, reducing 'server sprawl' and infrastructure costs, particularly for large data centres and managed service providers (MSPs) with cloud-based offerings. The new Linux on Power agent enhances Halcyon's Network Server Suite, which already includes a Linux on Intel agent, as an optional plugin module. Network Server Suite is a leading monitoring and automation solution that supports Windows, AIX and Linux operating systems. It ensures maximum availability of critical servers, applications, processes, services and network devices to maintain business operations around the clock.

## QueueMetrics Asterisk call-centre monitor version 14.10 is out

QueueMetrics version 14.10 is centred around three main themes -- improving usability by implementing a series of suggestions that came as feedback on the Icon page, incorporating a number of changes that make the experience generally smoother, and improving deployability by leveraging and extending the new JSON API. This version also has a command-line updater tool and a new HTTP-only data loader that does not require a direct connection to the QueueMetrics database. This makes deploying large numbers of cloud QueueMetrics instances easier. Plus, a number of configuration defaults have been updated and the sample

## Canonical may stop producing 32-bit Ubuntu by 2016

Canonical has no immediate plans of dropping the development of 32-bit Ubuntu ISOs. On being asked about when 32-bit images would stop being made, company spokespersons said there were no plans for it as of now, but these could be dropped after the release of Ubuntu 16.04. Ubuntu 16.04 LTS is supposed to be released in 2016, which should be the last 32-bit release of Ubuntu distro. There is a lot of ground work that Ubuntu developers need to do in order to push people to adopt 64-bit Ubuntu. They need to investigate 32-bit EFI issues and push the 64-bit Linux kernel to 32-bit Ubuntu users on x86\_64 systems. Ubuntu developers have started to push people to adopt 64-bit compatible hardware. Last year, Ubuntu started recommending the 64-bit version of Ubuntu as default to users. The latest Ubuntu 14.10 64-bit has a way better benchmark score than the 32-bit version.

## New kernel live patching solution out

SUSE launched its live kernel patching software, called kGraft, in February. Red Hat released kPatch in April. Even though the news created a big buzz amongst users, both have not been accepted at the upstream level yet. Now, a new live kernel patching solution based on ftrace mechanism is out. This is a combination of the best functionality in kpatch and kGraft. It uses the kernel interface with the kernel module function for live patching. This solution was released at a live patching mini conference at the Linux Plumbers Conference, which happened recently. The new live kernel patching mechanism can be found in the patch series released last month. Developers will be able to test it out when Linux kernel v3.19 is out.

## HostForLIFE.eu proudly announces Drupal 7.3 hosting

The ASP.NET spotlight hosting partner in Europe, HostForLIFE.eu, has announced the availability of new hosting plans that are optimised for the latest update of the Drupal 7.3 hosting technology. Drupal 7.3



is a free, open source Web development platform for online content and user communities. It is a free software package that allows customers to easily organise, manage and publish customer content. Drupal 7.3 is used by thousands of high profile websites and is subject to rigorous security testing, both by the Drupal community and by security experts around the world. It is a powerful solution that lets marketers and others build and manage sites, multi-channel digital experiences, shopping experiences, and integrate a diverse ecosystem of digital marketing and related tools and technology. Drupal is flexible, open, easy to integrate with, and easy to use for marketers and non-technical people - and it easily connects to marketing automation, CRM, email marketing, analytics, optimisation and other marketing technology solutions. HostForLIFE.eu is a popular online Windows-based hosting service provider. The company has managed to build a strong client base in a very short period of time. It is known for offering ultra-fast, fully-managed and secured services in the competitive market. Its powerful servers are especially optimised and ensure Drupal 7.3 performance. Drupal has perhaps the best data centres on three continents, unique account isolation for security, and 24/7 proactive uptime monitoring.

database includes a better sample QA form, fixing a number of bugs and issues that would annoy users. The most important is that call attempts are now filtered correctly. QueueMetrics 14.10 fully supports the upcoming Asterisk 13, which will be the core Asterisk LTS release coming out in the next months, as well as all other Asterisk versions. This release will initially come with Catalan localisation, and more languages are on their way.

## Open source has a major role in data centre transformation

Large data centres are dynamically scalable, automated and have virtualised platforms for efficient access. Big enterprises need to adopt open systems for their infrastructure, as independent applications and workloads are their top priority. There has been a massive increase in the use of open source systems by cloud and service-based IT



models. Big organisations such as the Singapore Stock Exchange, Wall Street firms, and some global investment banks use open source infrastructure. But open source models are yet to penetrate beyond particular infrastructure layers. There are some challenges that every

evolving technology goes through. If some structural challenges are overcome, open source infrastructure can penetrate the market in coming years.

A panel of leading IT experts met in Singapore recently to discuss the challenges faced by open source infrastructure. Enterprises are leaning more towards open source as they are adopting the cloud within their infrastructure. This is mainly because, irrespective of the technology platform, SaaS (Software-as-a-Service) enterprises are concerned only about the output from their technology. Another reason is that the cloud involves the use of a large number of software segments, which one software vendor or enterprise cannot excel in. By mixing and matching environments and parts of different builds, software platforms have become the key to innovative development. And open source enables and promotes such collaborative type of development. The huge community that Linux and OpenStack have built is adding endless possibilities of innovation every other day.

## Debian 9 and Debian 10 receive their code names

The names have been finalised and announced for the next two versions of the Debian operating system—Debian 9 and Debian 10. These two versions are



expected to arrive in the coming years. The code name for Debian 9 is Stretch and for Debian 10 is Buster. Earlier, no formal code names were ever announced for Debian releases, but it's time for change. These days, Debian users know the versions by their names mostly and not by numbers. For instance, users know the Wheezy One is quite stable, while Jessie is still in its testing phase. The Debian 8.0 branch has been in the pipeline for some months now and is due for final release soon. The Debian team is also working on unblocking and fixing bugs, as per requests, and sixteen bugs still need their close attention.



# Indian Government Opts for Open Source in New Policy!

Open source experts are quite in favour of the new policy and believe this is the right move on the government's part. However, they caution that it must be implemented with extra care for the best impact.

Open source technology is set to make a strong impact in India. An all-new policy on open source software is the latest addition in the Indian government's initiatives towards making the country more tech-savvy. Under this new policy, all government software applications are to be switched to open source, leading to huge savings and marking an end to the use of proprietary software.

## A part of the 'Digital India' initiative

The policy is being adopted as a part of the Indian government's 'Digital India' initiative, which will ensure more effective governance for the citizens of this country. This initiative will make government services accessible to citizens across the nation, digitally, which means the paperwork will be reduced gradually. Rural areas will also be connected through high-speed Internet, as a part of this project. The newly adopted open source policy will have a strong influence on this initiative.

## The rapid spread of e-governance in the country

E-governance is on in full swing across the country, and information and communication technology has become an integral part of government services. The aim is to build a better network between the common people of India and the Central as well as state governments. As Venkatesh Swaminathan, country head, The Attachmate Group India (Novell, SUSE, NetIQ, Attachmate), puts it, "State governments have already started lots of e-governance initiatives, and it has spread across services like land records, birth and death registrations, and even the provident fund."

## A GitHub-like repository for open sourcing code

With the open source policy coming into effect, the government will also create a GitHub-like repository for all open source software. This will allow developers to contribute to the projects in collaboration with each other, as the code of the software will be shared for all. If collaborative development is encouraged further, it will become easier to use open source software whenever and wherever required. More applications will be derived from this kind of collaborative development.

## A cost-effective measure

Millions of dollars are expected to be saved on the purchase of IT products and related services, once the open source policy is adopted effectively in the offices of the Indian government. The implementation of this open source policy will prove highly


cost-effective for the government as, whenever the government requires any software product, it will not need to pay for it any more. This will consequently have an impact on the savings strategy of the Indian government.

## Open source in the government sector is not a new concept

The use of open source in government software applications began long ago. Swaminathan says, "The government is also working with its internal agencies to develop open source offerings. It has taken up a lot of initiatives, specifically to develop applications on open source technology, to provide them on an enterprise-class scale along with the underlying infrastructure, as the software can be managed, patched and is secure." He also says that open source projects can be very successful across sectors like education, defence-related projects like DRDO, and more.

## What's so special about the new Open Source Policy?

According to the Department of Electronics and Information Technology (DeitY), the new policy will allow the source code of any software application to be modified by a user, as per individual requirements. The common repository will help in reducing duplication of applications too. If any government agency or department wants to use any particular application in its own way, then customisation of the code will become easier through this common repository. The new policy will also help to speed up the deployment of IT. Open source experts are quite in favour of the new policy as they believe this is a right move on the government's part. But they warn that extra care must be taken in the implementation of the policy, to ensure the best impact.

According to Swaminathan, the best qualities of open source technology are its flexibility and the wide choice it offers. Proprietary technologies demand complete dependence on the software vendor for all the solutions as well as developments, while open source technology is unquestionably open for all. He is also hopeful that the new policy, "...will offer the best benefits to people, whom the government is serving, and it will also be great for netizens like us who're able to access these services." Apart from creating an e-literate society and boosting savings, the policy is also aimed at accelerating the digital drive across the country.  **END**

## By: Sanchari Banerjee

The author is a member of the editorial team in the EFY Group. She loves to explore new innovations in the technology world.

# Microsoft Releases Source Code of .NET for Linux and Mac OS!

Microsoft is accumulating all the next-gen .NET components under a single umbrella, named .NET 2015.

It seems like the sky is the limit for Microsoft, as Satya Nadella leads it on a journey of rapid reinvention. It's a historical moment for Microsoft as its very own .NET framework has been open sourced completely, and that too to function across platforms like Linux and Apple Mac OS, the arch rivals of the Redmond tech giant's Windows OS. The crucial move towards broadening Microsoft's footprint in the world of open source computing was announced at Connect(), the virtual developers' event in New York City.

## .NET as open source, though not for the first time

It's not the first time that .NET has been open-sourced, though. In April this year, at the Build conference, the first step was taken to open source the platform, and multiple .NET components were released as open source projects, under the .NET Foundation. But now the process has been taken to the next stage of development, as the entire source code has been open sourced and the versions will be made available for diverse platforms including Linux and OS X. Now, .NET will pose tough competition for Oracle's Java, which is already a cross-platform software development solution. The latest move by Microsoft will attract developers who use a variety of open source technologies and build cloud applications through multiple components. The .NET framework will no longer remain the programming model for developers who build apps only on Windows.

## .NET: Not dependent on the Mono project

If we're to go by the recent development history of Microsoft, the latest announcement is not such a huge shock. According to S. Somasegar, corporate VP developer division, Microsoft Corp, the company now has an offering, which will be both relevant and attractive for developers working on any application. In this way, Microsoft is grabbing every possible opportunity to stay on top in the modern software and online services era, rather than resting on its laurels thanks to Windows reigning supreme on the desktop.

Earlier, the Mono project allowed developers to use the .NET framework on non-Windows platforms, but now that dependence is no longer necessary. The Mono project is an open source implementation of the .NET stack, and it's managed by Xamarin, a vendor of cross-platform mobile development tools. A new GitHub repository has also been created, where the code of .NET will be open sourced.


Somasegar is hoping that the switch to open source will be completed in the next few months.

There are several advantages of .NET as an open source package, which go much beyond just simple community-based development. Developers will be permitted to include .NET framework components in container applications as well. .NET components will be open sourced under the Apache 2.0 licence, and this will allow developers to use the code in proprietary projects too.

## The history of .NET and Microsoft's struggles in the open source market

.NET dates back to 2002 when it was released by Microsoft for developers to build applications on Windows. During its development stage, several internal complications arose within the company. Then, experimentation started with cross-platforms in the form of the Silverlight framework, which was an implementation of .NET to run on the Mac OS. But support for Silverlight didn't continue for long. Then, during 2007-08, open source technologies started emerging fast in the marketplace, as Google released the Android OS as open source software and GitHub opened its repositories for open source code. Though the world was changing fast, Microsoft was unable to keep up with the advancements.

## .NET today

Ever since Satya Nadella took charge of the company as its CEO, things have changed for Microsoft. The company is now striving hard to meet developers' requirements. Microsoft is accumulating all the next-gen .NET components under a single umbrella, named .NET 2015. This also includes .NET 4.6, which is the next update to the desktop .NET framework. Somasegar believes this new development will open up a bunch of new opportunities for Microsoft's partners as well. He acknowledges that Nadella has played a crucial role in the direction Microsoft is going with regard to its openness, and commends Nadella for his foresight and continuous drive towards more and more development. Certainly, an open source .NET will make Microsoft software more appealing to developers.  **END**

**By: Sanchari Banerjee**

The author is a member of the editorial team in the EFY Group. She loves to explore new innovations in the technology world.



# OSFY Classifieds

Classifieds for Linux & Open Source IT Training Institutes



## WESTERN REGION

### Linux Lab (empowering linux mastery)

**Courses Offered:** Enterprise Linux & VMware

**Address (HQ):** 1104, D' Gold House, Nr. Bharat Petrol Pump, Ghyaneswari Paduka Chowk, FC Road, Shivajinagar Pune-411 005

**Contact Person:** Mr. Bhavesh M. Nayani

**Contact No.:** +020 60602277, +91 8793342945

**Email:** info@linuxlab.org.in

**Branch(es):** coming soon

**Website:** www.linuxlab.org.in

### Linux Training & Certification

**Courses Offered:** RHCSA, RHCE, RHCVA, RHCSS, NCLA, NCLP, Linux Basics, Shell Scripting, (Coming soon) MySQL

**Address (HQ):** 104B Instant Plaza, Behind Nagrik Stores, Near Ashok Cinema, Thane Station West - 400601, Maharashtra, India

**Contact Person:** Ms. Swati Farde

**Contact No.:** +91-22-25379116 / +91-9869502832

**Email:** mail@lfcert.com

**Website:** www.linuxlab.org.in

## NORTHERN REGION

### GRRAS Linux Training and Development Center

**Courses Offered:** RHCE, RHCSS, RHCVA, CCNA, PHP, Shell Scripting (online training is also available)

**Address (HQ):** GRRAS Linux Training and Development Center, 219, Himmat Nagar, Behind Kiran Sweets, Gopalpura Turn, Tonk Road, Jaipur, Rajasthan, India

**Contact Person:** Mr. Akhilesh Jain

**Contact No.:** +91-141-3136868 / +91-9983340133, 9785598711, 9887789124

**Email:** info@grras.com

**Branch(es):** Nagpur, Pune

**Website:** www.grras.org, www.grras.com

## SOUTHERN REGION

### Advantage Pro

**Courses Offered:** RHCSS, RHCVA, RHCE, PHP, Perl, Python, Ruby, Ajax, A prominent player in Open Source Technology

**Address (HQ):** 1 & 2, 4th Floor, Jhaver Plaza, 1A Nungambakkam High Road, Chennai - 600 034, India

**Contact Person:** Ms. Rema

**Contact No.:** +91-9840982185

**Email:** enquiry@vectratech.in

**Website(s):** www.vectratech.in

### Duestor Technologies

**Courses Offered:** Solaris, AIX, RHEL, HP UX, SAN Administration (Netapp, EMC, HDS, HP), Virtualisation (VMWare, Citrix, OVM), Cloud Computing, Enterprise Middleware.

**Address (HQ):** 2-88, 1st floor, Sai Nagar Colony, Chaitanyapuri, Hyderabad - 060

**Contact Person:** Mr. Amit

**Contact No.:** +91-9030450039, +91-9030450397.

**Email:** info@duetor.com

**Website(s):** www.duestor.com

### IPSR Solutions Ltd.

**Courses Offered:** RHCE, RHCVA, RHCSS, RHCD, RHCA, Produced Highest number of Red Hat professionals in the world

**Address (HQ):** Merchant's Association Building, M.L. Road, Kottayam - 686001, Kerala, India

**Contact Person:** Benila Mendus

**Contact No.:** +91-9447294635

**Email:** training@camaticindia.com

**Branch(es):** Kochi, Kozhikode, Thrissur, Trivandrum

**Website:** www.ipser.org

### Linux Learning Centre

**Courses Offered:** Linux OS Admin & Security Courses for Migration, Courses for Developers, RHCE, RHCVA, RHCSS, NCLP

**Address (HQ):** 635, 6th Main Road, Hanumanthnagar, Bangalore - 560 019, India

**Contact Person:** Mr. Ramesh Kumar

**Contact No.:** +91-80-22428538, 26780762, 65680048 / +91-9845057731, 9449857731

**Branch(es):** Bangalore

**Email:** info@linuxlearningcentre.com

**Website:** www.linuxlearningcentre.com

## Eastern Region

### Academy of Engineering and Management (AEM)

**Courses Offered:** RHCE, RHCVA, RHCSS, Clustering & Storage, Advanced Linux, Shell Scripting, CCNA, MCITP, A+, N+

**Address (HQ):** North Kolkata, 2/80 Dum Dum Road, Near Dum Dum Metro Station, 1st & 2nd Floor, Kolkata - 700074

**Contact Person:** Mr. Tuhin Sinha

**Contact No.:** +91-9830075018, 9830051236

**Email:** sinhatuhin1@gmail.com

**Branch(es):** North & South Kolkata

**Website:** www.aemk.org

**Want to  
Advertise in  
This Section?**

**Call Omar On  
+91-995 888 1862**

**or**

**write to omar.farooq@efy.in**

# How to Select the Right Desktop Scanners

It's a challenge to find the right scanner. You cannot just decide on any scanner as these vary in shape, size and also the purpose they serve. Here is a guide to make sure you pick the right scanner to suit your requirements.

**D**esktop scanners are specialised devices designed to make speedy work of removing piles of paper from the desk and transferring all that content onto the desktop or computer. This is the product that does the job of managing files and documents quickly and accurately. Gone are the days when only Internet cafes had scanners. If you still keep a bundle or albums of photographs, or a file full of hard copies of important documents, a desktop scanner can make your life easy and manageable. A scanner can not only bring your memories back to life whenever you want to view your photograph collection, but it also ensures both pictures and documents are safely stored and easy to manage. "Today, there are a lot of documents at home, like loan documents, insurance policies, personal certificates, old school photographs which are printed, as earlier there were no digital photographs...and we would like to keep all these documents secure with us. So we have introduced a product keeping all these things in mind, which can be carried anywhere and weighs just 550 grams; you can connect it to a laptop without using any external power (USB)," explains Jaspreet Kohli, senior manager, sales and marketing, IMS, Canon India Pvt Ltd. Desktop scanners are also great at managing heavy or bulky documents like magazines or books. Most scanners these days connect to the computer or PC through USB cables, but some are even operated wirelessly.

The demand for scanners is increasing rapidly. Today's scanners are highly specialised, and they vary in shape, size, price and even the purpose they serve. The market offers scanners of different specifications, to address varied requirements. According to Kohli, "We go to different kinds of customers, try to understand their requirements and then try to bring in that technology or the product required for the Indian market. Our USP is not just the technology but our service. We cater to more than 3000 cities in India. Today, what customers look at is also the after sales service, and Canon definitely offers good service to the customers. We also give end-to-end solutions."

## Factors to be considered while buying a desktop scanner

### Portability

Portability is one of the most important factors a buyer should consider while buying a scanner. From the

innumerable scanners in the market, you must first look for the brands that offer portable scanners instead of the heavy, bulky ones, and then compare other aspects. The target audience of portable scanners comprises the home segments, SOHO, or professionals like lawyers who cannot carry huge files everyday. Having said that, many portable models are fairly basic, which end up scanning only one side of the paper at one time. So, for office use, you may not want a portable scanner.

### Resolution

You would not want a document or a photograph scanned with a low resolution. So another factor to keep in mind while selecting a scanner is the resolution it offers. But you will have to be very careful about keeping your documents clean, as most high resolution devices will magnify dust or scratches on the document even more clearly.

### Flatbed or sheet feeder?

Select a flatbed scanner in case you don't require to scan three dimensional objects or documents longer than one or two pages. Opening the flatbed and adjusting the sheet once or twice a day will not be a problem, but having to repeat this process 10 times a day can be tedious and that is when you need a sheet feeder. So choose between a flatbed scanner and a sheet feeder keeping your requirements in mind.

### Duplex

A duplex scanner is one that scans both the sides of the document at the same time. These models have dual scan elements, enabling them to scan on both the sides simultaneously. If you do not require this feature too often, then a manual duplexing in the driver is the most economical choice.

### Software

Almost all the scanners available in the market work with just about any scanning related software, but if particular software that you need already comes with the scanner, you would not have to pay extra for it. Some of the software these days offer features like photo editing, text indexing, optical character recognition (OCR), a business card program or the ability to create PDF documents.



## Some of the best scanners to select from



### Canon imageRUNNER 2002N

- **Scan speed:** 20ppm
- **Scan resolution:** 300×300 dpi
- **Processor:** 400MHz
- **Weight:** 35.5 kg

### HP Scanjet 200

- **Scan speed:** 10 x 15 cm (4 x 6 in) colour photo to file (200 dpi, 24-bit, about 21 sec for single image scan' after tiff);
- **Scan resolution:** 2400×4800 dpi
- **Processor:** 30MHz
- **Weight:** 20.86 kg



### Epson GT-S55

- **Scan speed:** 25ppm
- **Scan resolution:** 600 dpi
- **Weight:** 4.9 kg

### Brother ADS-1600W

- **Scan speed:** 18ppm
- **Scan resolution:** Up to 1200×1200 dpi (enhanced)
- **Processor:** ARM946 288MHz
- **Weight:** 1.6 kg



**Errata:** On Page 22 of the November 2014 issue of *Open Source For You*, the heading in the Buyers' Guide carried an error. The correct heading is: 'Some of the best network switches available in the Indian market'. The error is regretted.

**By: Manvi Saxena**

The author is a part of the editorial team at EFY.





With **7500+**  
registrations, we lived  
up to our tag of  
**Asia's #1  
Open Source  
Conference**





11th Edition

# OPEN SOURCE INDIA

NIMHANS Convention Center

7 - 8  
November  
2014

BENGALURU

# Thank You

**VISITORS**, for making the  
event a huge success

**SPEAKERS**, for contributing  
your valuable thoughts

**PARTNERS**, for your support



## See you again in 2015

[www.osidays.com](http://www.osidays.com)

For more details, call Omar (+91-9958881862) or email us at [info@osidays.com](mailto:info@osidays.com)

# CODE SPORT



In this month's column, we feature a set of interview questions on algorithms, data structures, operating systems and computer architecture.

**F**or the past few months, we have been discussing information retrieval, Natural Language Processing (NLP) and the algorithms associated with them. In this month's column, we focus on an important aspect of NLP known as textual entailment.

## What is textual entailment?

Let us understand more about this by first looking at the two following snippets of text:

- (1) India has a number of nuclear power plants which are used to generate electrical power. The first nuclear power plant was started in Tarapur, Maharashtra. The latest one was commissioned at Kudankulam, Tamilnadu.
  - (2) Kudankulam in Tamilnadu has a power generation station, which generates electricity.
- Now the NLP problem posed is to determine whether Snippet 2 can be inferred from Snippet 1. When human beings parse the text of Snippet 1 and Snippet 2, it is very easy for us to determine whether the latter can be inferred from the former. On the other hand, it is not easy for an automated NLP algorithm to reach this conclusion. This is the problem area that textual entailment techniques attempt to solve. Formally, given a text snippet 'T' and a text snippet representing the hypothesis 'H', a textual entailment program could determine whether they formed a textual entailment pair. 'T' and 'H' form a textual entailment pair, if a human reading 'T' and 'H' would be able to infer 'H' from 'T'.

Consider the following example of two snippets:

- (3) The Director of Public Health said, "It is important to stress that this is not a confirmed case of Ebola."

- (4) A case of Ebola was confirmed in Mumbai.

Now the question is whether Snippets 3 and 4 constitute a textual entailment pair? The answer of course is obvious to humans—that they do not form a textual entailment pair. But as we will see later, this is not a simple case for an automated textual entailment technique. Many techniques that use surface level string parsing or that are keyword-based, do wrongly mark this pair as textually entailed. Much of the complexity of automatic textual entailment techniques is needed to deal with and avoid such false positives.

An area closely related to textual entailment is paraphrasing. Two statements are paraphrases if they convey almost the same information. Consider the following snippets:

- (5) Ayn Rand wrote the book 'Fountainhead'.
- (6) Fountainhead was written by Ayn Rand.
- (7) Ayn Rand is the author of Fountainhead.
- (8) Fountainhead is the work of Ayn Rand.

Now statements 5, 6, 7 and 8 are all paraphrases of each other. If two statements A and B are paraphrases, then they are mutually textually entailed. In other words, both (A, B) and (B, A) are textual entailment pairs. Statement A can be inferred from B and vice versa. But textual entailment does not automatically imply that they are paraphrases. Consider our previous example: Statement 1 implies Statement 2, but not the other way around.

Textual entailment programs can be used for: (a) recognising whether two pairs of natural language expressions form a textual entailment pair, (b) given a single natural language expression, generate all possible TE expressions for it, and (c) given a document or set of documents, extract



all TE pairs. Similarly, paraphrase programs can be either recognisers, generators or extractors. Before we look into the techniques for TE and paraphrasing recognition/generation/extraction, let us look at the practical applications for TE and paraphrasing in NLP applications.

## **Applications of paraphrasing and textual entailment**

So far, we have been considering textual statements as inputs to paraphrasing systems. But this may not always be the case. In fact, one of the earliest NLP applications of paraphrasing was in the field of automatic ‘Question Answering’ (QA) systems. Consider a QA system in which the system is given a set of documents, and needs to find the answer to the posed question from among the documents. Given that the answer, if it is indeed present in the document, may be phrased differently from the way the question has been framed, it may be necessary to generate paraphrases of the question and check if any of the paraphrased questions can be answered using the content in the document collection. Let us consider a simple example in which the question being posed to the QA system is, “Who is the author of the book ‘Crime and Punishment’?” and one of the documents in the collection contains the passage, “Some of Leo Tolstoy’s finest works include ‘Crime and Punishment’ and ‘Anna Karenina’.” In this case, by paraphrasing the question as, “Whose work is ‘Crime and Punishment’?” the QA system may be able to easily figure out the answer as ‘Leo Tolstoy’. Instead of paraphrasing the question, in some cases, the possible answers may also be paraphrased to check if any of the paraphrases can serve as an answer to the question.

Text summarisation is another area where paraphrasing techniques are used widely. Given a set of text documents to be summarised, one of the important functions of a summary extractor is to identify the most important sentences from the texts to be summarised. For example, let us consider that the task given is to create a summary from all news articles in the last one month which discuss the Ebola virus epidemic. Since these newspaper articles are discussing the same event, many of the documents will contain sentences that convey almost the same information. Hence, it is important to avoid selecting sentences that are paraphrases in the summary. Hence, paraphrasing techniques can be applied to a particular sentence to check if it is a paraphrase of any of the existing sentences in the summary, and if found to be so, it can be discarded. Similar to paraphrasing, TE

can also be applied on a particular statement to check if it can be inferred from any of the existing sentences in the summary. And if so, then the statement can be discarded. Note that paraphrasing can also be used to achieve sentence compression, since it can help to generate a sentence which is shorter than the original sentence but conveys the same information.

Machine translation is another major area where paraphrasing and textual entailment techniques are applied. Paraphrasing, in particular, has been widely used in training machine translation systems by using a human generated translation reference and checking to see if the machine generated translation is a paraphrase of the human generated one. This is typically useful when popular literary works, which have been translated by humans, are used to train the machine translation systems. Let us consider a simple example, where the two statements are:

(a) Company A filed a case against Company B for copyright infringement.

(b) Company A accused Company B for copyright violation.

If our machine translation system has never encountered the phrase ‘filed a case’ in the source language during its training samples, it could try finding a translation for the paraphrased source sentence (b), if it has encountered the word ‘accused’ before in its training samples. Using paraphrasing allows a machine translation system to deal with words that it has not encountered before in its training samples.

Other areas where TE and paraphrasing techniques are useful include Information Extraction (IE) systems, which use automated techniques to extract information on specified topics from various sources. These are typically used for answering natural language queries in information retrieval systems. Consider, for example, a system that extracts information on all motor vehicle accidents from newspaper reports. Consider the following statement:

(a) The tourist bus collided with an oncoming car and crashed on to the median.

If this can be paraphrased as “There was an accident involving a bus and car,” then it is possible for the IE system to include this as a candidate news item in its search. Other areas in which TE and paraphrasing can be applied include automatic grading of students’ answers, and in providing simplified summaries of the expert documents in a form understandable to laymen.

We will continue our discussion on TE and paraphrasing next month, when we will look at the algorithms for them. Meanwhile, I would like to give readers the following assignment.

*...Continued on page 38*



Anil Seth

# Exploring Big Data on a Desktop, with Hadoop, Elasticsearch and Pig

In continuation of earlier articles, the author goes further into the subject to discuss Elasticsearch and Pig, and explain how they can be used to create an index for a large number of PDF files.

If the files and data are already in Hadoop HDFS, is Elasticsearch still useful? How does one create an index?

Consider a large number of PDF files that need to be searched. As a first step, process each PDF file and store it as a record in an HDFS file. Then, you may experiment with two different but very simple approaches to create an index.

- Write a simple Python mapper using MapReduce streaming to create an index.
- Install the Elasticsearch-Hadoop plugin and create an index using a Pig script.

The environment for these experiments will be the same as in the earlier articles – three virtual machines, *h-mstr*, *h-slv1* and *h-slv2*, each running HDFS and Elasticsearch services.

## Loading PDF files into Hadoop HDFS

Enter the following code in *load\_pdf\_files.py*. Each PDF file is converted to a single line of text. Any tab characters are filtered so that there are no ambiguities when using a Pig script. For each file, the output will be the path, tab, file name and the text content of the file.

```
#!/usr/bin/python
from __future__ import print_function
import sys
import os
import subprocess
# Call pdftotext to convert the pdf file and store the result
in /tmp/pdf.txt
def pdf_to_text(inpath,infile):
    exit_code=subprocess.call(['pdftotext','/'.
join([inpath,infile]),'/tmp/pdf.txt'],stderr=ErrFile)
    return exit_code,'/tmp/pdf.txt'
# Join all the lines of the converted pdf file into a single
string
# Replace any tabs in the converted documents
# Write the file as a single line prefixing it with the path
and the name
def process_file(p,f):
    exit_code,textfile = pdf_to_text(p,f)
    if exit_code == 0:
        print("%s\t%s"%(p,f), end='\t')
        print("%s" % ' '.join([line.strip().replace('\t',' ')
for line in open(textfile)]))
```

```
# Generator for yielding pdf files
def get_documents(path):
    for curr_path,dirs,files in os.walk(path):
        for f in files:
            try:
                if f.rsplit('.',1)[1].lower() == 'pdf'
                    yield curr_path,f
            except:
                pass
# Start here
# Search for each file in the current path of type 'pdf' and
process it
try:
    path=sys.argv[1]
except IndexError:
    path='.'
# Use an error file for stderr to prevent these messages going
to hadoop streaming
ErrFile = open('/tmp/err.txt','w')
for p,f in get_documents(path):
    process_file(p,f)
```

Now, you can run the above program on your desktop and load data into a file in Hadoop HDFS as follows:

```
$ ./load_pdf_files.py ~/Documents [HADOOP_USER_NAME=fedora \
hdfs dfs -fs hdfs://h-mstr/ -put - document_files.txt
```

## Using MapReduce to create an index

Log into *h-mstr* as user *fedora* and enter the following code in *'indexing\_mapper.py'*.

```
#!/usr/bin/python
import sys
from elasticsearch import Elasticsearch
# Generator for yielding each line split into path, file name
and the text content
def hdfs_input(sep='\t'):
    for line in sys.stdin:
        path,name,text=line[:-1].split(sep)
        yield path,name,text
# Create an index pdfdocs with fields path, title and text.
# Index each line received from Hadoop streaming
```



```
def main():
    es = Elasticsearch(hosts='h-mstr')
    for path,name,text in hdfs_input():
        doc = {'path':path,'title':name, 'text':text}
        es.index(index='pdfdocs', doc_type='text', body= doc)
if __name__ == "__main__" :
    main()
```

Run the code in the following command on *h-mstr*:

```
$ hadoop jar /usr/share/java/hadoop/hadoop-streaming.jar \
-files indexing_mapper.py -mapper indexing_mapper.py \
-input document_files.txt -output es.out
```

The following URLs will give you more information about the allocation and status of the index.

```
http://h-mstr:9200/_cat/allocation
http://h-mstr:9200/_cat/indices
```

## Using a Pig script to create an index

The Fedora 20 repositories do not as yet have the Pig distribution. It will be included in Fedora 21. So, download and install Pig from the Apache site <http://pig.apache.org/releases.html> on each of the virtual machines.

You will also need to install the Elasticsearch-Hadoop plugin on these systems. For example, you may run the following commands from *h-mstr*:

```
$ sudo /usr/share/elasticsearch/bin/plugin \
-u http://download.elasticsearch.org/hadoop/elasticsearch-
hadoop-2.1.0.Beta2.zip \
-i elasticsearch-hadoop

$ ssh -t fedora@h-slv1 sudo /usr/share/elasticsearch/bin/
plugin \
-u http://download.elasticsearch.org/hadoop/elasticsearch-
hadoop-2.1.0.Beta2.zip \
-i elasticsearch-hadoop

$ ssh -t fedora@h-slv2 sudo /usr/share/elasticsearch/bin/
plugin \
-u http://download.elasticsearch.org/hadoop/elasticsearch-
hadoop-2.1.0.Beta2.zip \
-i elasticsearch-hadoop
```

The Pig script, *indexing.pig*, for creating the index is just four lines. The *elasticsearch-hadoop jar* file has to be registered. The Hadoop text file is loaded. The tuple (path, text, title) is stored in (indexed by) Elasticsearch storage.

```
REGISTER /usr/share/elasticsearch/plugins/hadoop/dist/
elasticsearch-hadoop-2.1.0.Beta2.jar;
A = LOAD 'document_files.txt' USING PigStorage() AS
(path:chararray, title:chararray, text:chararray);
B = FOREACH A GENERATE path, text, title ;
STORE B INTO 'docs/text' USING org.elasticsearch.hadoop.pig.
EsStorage();
```

You can check the status of the indices and compare the *pdfdocs* index created earlier with the *docs* index created by running the Pig script:


```
[fedora@h-mstr ~]$ pig indexing.pig
```

The ultimate test is to compare the results of the two indices; e.g., you can browse the Elasticsearch index searching for 'Python' in the contents and displaying up to 25 values. Only the values of the fields path and title will be displayed.

```
http://h-mstr:9200/docs/_search?pretty=true&size=25&fields=pat
h,title&q=text:python
```

The more flexible option is to use a *json* string to query as follows (for details, go to <http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/query-dsl.html>):

```
curl "h-mstr:9200/pdfdocs/_search?pretty=true" -d '{
  "fields": ["path", "title"],
  "size": 25,
  "query": { "query_string": { "query": "python" } }'
```

If all has gone well, you should get the same answers for the queries—whether you use the *docs* or *pdfdocs* indices. **END** 

### By: Anil Seth

The author has earned the right to do what interests him. You can find him online at <http://sethanil.com>, <http://sethanil.blogspot.com>, and reach him via email at [anil@sethanil.com](mailto:anil@sethanil.com)

◎ TechnoMail  
Enterprise Email Solutions

◎ TechnoMail Cloud  
Zero Capital Email Solution

**TechnoInfotech®**

info@technoinfotech.com  
www.technoinfotech.com

◎ High Volume SMTP Solutions  
(Hosted and on-Premise)

◎ Business Promotion Mailing  
(Hosted and on-Premise)

# Open Source India 2014

## The Industry Meets the Community

We take you on a tour of the 11th edition of Asia's largest convention on open source, featuring over 45 speakers in various sessions, hands-on workshops and success stories... a post show report by **Manvi Saxena**, a member of the editorial team at EFY, on an event that was filled to the rafters!

The banner is a vertical rectangle with a blue header containing the word 'WELCOME' in white. Below this, the text '11th Edition' is in a smaller font. The main title 'OPEN SOURCE INDIA' is in large, bold, black letters, with the 'O' in 'OPEN' being a stylized circle with green and orange segments. To the right of the title, the dates '7 - 8 November 2014' are listed. Below the title, the venue 'NIMHANS Convention Center' and city 'BENGALURU' are mentioned. A line of text below that reads 'Asia's #1 Open Source Convention'. The bottom section is titled 'POWERED BY' and contains several logos categorized by partnership level: Platinum Partner (HP Helion), Gold Partners (Microsoft, Oracle Linux, MongoDB, MySQL, Zimbra), Registration Partner (Wipro), Knowledge Partner (UNOTECH), Organisers (FOSD, EFYGROUP, OpenSource.com, EFYtimes.com, electronics), and Media Partner (OpenSourceForYou).

Open Source India 2014, a two-day event focusing on the collaboration between the industry and the community, witnessed a whopping 7500+ registrations. At the event, the queues outside registration booths began to form almost an hour before the scheduled time. People who already had a Silver Pass wanted to upgrade it to a

Silver Plus, Gold or Platinum, in order to attend more sessions and workshops. By noon of the first day, i.e., on November 7, the organisers reluctantly stopped issuing passes for several tracks and workshops because of non-availability of seats.

Taking open source technology to another level, OSI was packed with nearly 45 sessions by eminent speakers like Lux Rao (CTO, technology services, HP), C N Krishnan (programme director, Anna University), Aahit Gaba (open source attorney, Lyra Infosystems Pvt Ltd) and Jacob Singh (regional director, Acquia India). Talks on FOSS for everyone, the cloud, mobile app development, databases, Web development, kernel development, the OpenStack mini conference, IT infrastructure management and the 'Success Stories' track were all conducted on time.

### Major highlights of the event

**Record-breaking attendance for the keynotes and workshops:** Extra seating had to be arranged for all the keynote sessions and the workshops. All the four keynote sessions witnessed a major turnout and the crowds seemed to be a highly satisfied lot. The same could be said of the workshops, which were also a total hit and were much talked about.

**The 'Success Stories' track:** This was introduced in the 10th edition of Open Source India and was much talked about at the event by the speakers and the attendees. CIOs and CTOs of prominent companies like Rediff.com, Goibibo, bookmyshow.com, Justdial, Patterbuzz, Mahindra & Mahindra Financial Services Ltd, and Karunya University shared their experiences of implementing open source in their respective organisations. Talking to our correspondent after his talk, Sandipan Chattopadhyay, CTO, Justdial, said, "Well, it keeps getting more and more interesting every time and the crowd was more interactive. This time, there are a lot of young people who have come in and they are asking a lot of questions. The event is better organised, and is a lot more focused."

**The HP and Microsoft tech challenges:** Long

queues were witnessed at the HP Helion and Microsoft booths in between the sessions. There were many intriguing activities conducted at these two booths, with a lot of goodies being given away to the participants. Speaking about this, Rajiv Pandey, distinguished technologist, HP Cloud, said, "I think the event is great and it is amazing to see what's going on with open source here in India. I loved the energy and excitement here; it's really nice. While talking to people coming to the HP booth, I realised that they know so much about the new technologies and that they are very up-to-date, which was very gratifying. It's been a great event. It's already been 11 years of OSI and there clearly has been good work put into it. I would love to be a part of it in the years to come."

**Wipro Talent Hunt:** IT giant, Wipro Technologies, had tied up with Open Source India 2014 to recruit candidates for diverse profiles and a few applicants were shortlisted too.

## Day 1

The first day of the convention began with the welcome speech by Rahul Chopra, managing director, EFY Group. He spoke about Open Source India's 11-year journey and thanked everyone for the love they have shown for this event.

The sessions started with a keynote by Rajeev Pandey, distinguished technologist, HP Cloud, who spoke about 'Deployment architecture for OpenStack in the enterprise', which set the ball rolling for an eventful day for the IT professionals. He kept the audience enthralled with his vast knowledge and brilliant presentation skills.

The crowd then dispersed into the various halls and workshop rooms where different sessions were being conducted. The first half of Day 1 had 'FOSS for everyone', 'Cloud and Mobile App Development' tracks, and the speakers included Anthony Rees (HP Cloud), Jose Miguel Parrella (CTO, M&O, Microsoft Corp) and Michael Meskes (CEO, Credativ International GmbH), amongst others. Apart from these sessions, there were hands-on workshops going on simultaneously on topics like 'MySQL performance tuning', 'Drupal' and 'Architecting and building solutions using MongoDB'.

The second half of Day 1 also started with a keynote, this time by Ramesh Srinivasan, senior director, Linux and virtualisation, Oracle and Sanjay Manwani, director, MySQL, Oracle India. They enlightened the audience on how far open

source has travelled, discussing the state of the Penguin, and the Dolphin. Other post-lunch tracks included a mini conference on 'OpenStack and the kernel' as well as a workshop on 'Big Data'.

Everyone, including the exhibitors, speakers as well as the attendees, seemed satisfied with the efforts put in by the OSI content team and expectations grew for the second day.

## Day 2

With Day 1 being a great success, the second day was filled with open source enthusiasts eager to acquire more knowledge about the latest trends and technologies. Day 2 commenced with a keynote session by Jose Miguel Parrella, CTO, M&O, Microsoft Corp, who spoke about 'Innovations in the open source world', which was then followed by various tracks. Sessions on Day 2 included 'Success stories', 'Databases', 'Web app development' and 'IT infrastructure management' as well as workshops on 'Cross-platform mobile app development', 'Programming OpenStack', 'Software architecture in 90 minutes' and 'HP Helion OpenStack'.

The post-lunch keynote by Kamal Brar, VP, Asia-Pacific, MongoDB, was another hit and was dedicated to the innovations in the FOSS world. The OSI team went around getting as much feedback as possible, and was pleased to find out that everyone was satisfied with the show. Vikram Bhatia, director and technical evangelist, HP Cloud; Ryusuke Kajiyama, MySQL sales consulting manager, APAC; and Jose Miguel Parrella, Microsoft Corp, were interviewed by the OSI team on the second day and they spoke about their careers, their visit to OSI, their ventures and much more. Day 2 as well as the 11th edition of Open Source India ended well, with the hope of taking this event to an even higher level. All the speakers, exhibitors, attendees and the sponsors have extended their support for the 12th edition of OSI.

The NIMHANS Convention Center is huge, and to fill it up with the right people is a challenge in itself. The stalls of the partners, which included HP Helion, Microsoft, MongoDB, Oracle, Zimbra, Wipro and Unotech, were beautifully set up. The speakers were very enthusiastic about sharing their feedback with us. Rajesh Sola (from C-DAC's advanced computing training school, Pune), who spoke on 'Android Open Accessory Protocol-Turn your Linux Machine into an Android Development Kit (ADK)', said, "I have enjoyed the event thoroughly, and I just love the idea of enterprises connecting with technology. I look forward to next year."



## OSI Days 2014 Fact Sheet

- Footfalls: 1500 unique visitors
- No. of speakers: 60+
- No. of hands-on workshops 8
- Partners: Microsoft, HP Helion, Zimbra, Unotech, Wipro, MongoDB, and Oracle
- Registrations: 7500+



I go to many open source events all around the world, where the industry meets innovation. Open source is the core of such innovation. I am really excited to be here in India. For Microsoft, this is the fifth year for supporting Open Source India. Looking at the community, the opportunity is really big -- the strength of the community is really big. The cloud is really changing and so is the technology. OpenCloud, OpenStack and open source are the future of our industry.

We're committed to supporting this community and its projects, and to our partners and customers using open source.

### *Any suggestions for the next year...*

Well, I think there are a number of projects being represented in this event. OpenStack is one of them, but there are many other technologies out there, such as containers, Docker, etc, which compose the ecosystem of open source. So I think these should be the topics to be covered next year. I would definitely love to be associated with OSI the next year.

—**Jose Miguel Parrella, CTO, M & O, Microsoft Corp.**



OSI was more interesting this time and there was more interaction with the crowd. This wasn't the case last year; there were knowledgeable people last year too, but this time there were a lot of young people who asked a lot of questions. The event was better organised and a lot more focused. It was a plus plus event, as compared to the last year.

### *Any suggestions for the next year...*

Unfortunately, I couldn't spend much time at the event, but I do believe that the workshops have increased, which is a good thing. Maybe you could end the event with a hackathon or have a place for start-ups to come and talk about their problems. You should plan the program for the next year, too, and share it at the event. You could announce the big activities you plan to conduct the next year, so that people can plan their participation for the next year too. I look forward to OSI 2015..

—**Sandipan Chattopadhyay, CTO, Justdial**



I think the event is great and it is amazing to see what's going on with open source here in India. I loved the energy and excitement here. The people who came to the HP booth knew so much about the new technologies and were very up to date with respect to the knowledge on open source. It's been a great event. Maybe we could have more tracks with more specialities, and have the event for one more day. The event has been conducted already for 11 years, and

there is clearly good work put into it. I would love to be a part of it next year too..

—**Rajiv Pandey, distinguished technologist, HP Cloud**



I was really surprised, positively surprised by this event. I was kind of expecting a smaller event but this is really, really huge. A lot of people are so enthusiastic about open source. I think one can learn a lot here; a lot of people were really looking forward to learn on this platform, which is very positive.

### *Any suggestions for the next year...*

There are no changes I would want in the next one! Take it as a compliment.

—**Michael Meskes, CEO, Credativ International GmbH**



The Penguin was omnipresent at Open Source India, as always



A big thank you to all our advisors and speakers for making the event a grand success





FOSS enthusiasts make a beeline at the registration counter



Thank you for such an overwhelming turnout and making the show a success



For Tushar Deshpande from Zimbra, OSI 2014 translated into some real business



MongoDB connected with its target audience and created awareness about its platform



Oracle, a regular OSI supporter, influenced Linux lovers with the best of its technologies



HP Helion made a grand debut at OSI 2014 and saw a lot of traction from the audience too...



At OSI 2014, Microsoft demonstrated its love for Linux in some really innovative ways!



Record-breaking footfall at the keynote conducted by Rajiv Pandey



Microsoft's love for Linux was pretty visible at OSI 2014



Unotech Software Pvt Ltd, a valuable sponsor of OSI, creating awareness about its services



Wipro's recruitment drive at Open Source India was a major hit amongst the attendees



The developers had a great interaction and had a lot to take away from the experts at HP Helion



# Open Source India 2014



Sanjay Manwani, director, Oracle India, shared insights on Oracle's open source journey during his session



Ramesh Srinivasan, senior director, Oracle Linux & Virtualisation, conducted a keynote session on the 'State of the Penguin and the Dolphin'



Jose Miguel Parrella, CTO-M&O, Microsoft Corp, talked of how Microsoft endorses 'openness' in his keynote address



Rajiv Pandey, distinguished technologist, HP Cloud, gave a keynote session on 'Deployment Architecture for OpenStack in the Enterprise'



Kamal Brar, VP, APAC MongoDB, gave a keynote session on the path to innovation



Piyush Mathur, senior solutions advisor, Zimbra, talked about secure collaboration for the cloud



Vikram Bhatia, director, technical evangelism for HP Cloud, compared enterprise cloud platforms and technologies in his talk



Ryusuke Kajiyama, MySQL sales consulting manager, JAPAC, talked of the ways in which management efficiency and cloud sales could be improved for MySQL



Pooja Maheshwari, technical architect, Transilcity, showed the road to a mature mobile enterprise



Jacob Singh, regional director, Acquia Inc, talked about personalisation on the Web, the Drupal Way



Michael Meskes, CEO, Credativ International GmbH, conducted an interesting session on how the open source community could meet with businesses



Aahit Gaba, open source attorney, Lyra Infosystems Pvt Ltd, shared information on how to license the mobile app



Sandipan Chattopadhyay, CTO, Justdial, shared the very interesting journey of open source at Justdial



Sanil Kumar D, from Huawei Technologies India, gave a detailed talk on the Input/Output Memory Management Unit (IOMMU)



Anthony Rees, HP Cloud, talked on transforming businesses with DevOps and open hybrid clouds



Krishna Kumar, architect, HP, talked about Trafodion - an enterprise-class SQL based on the Hadoop DBMS engine





Rajesh Sola from CDAC's Advanced Computing Training School, Pune, enlightened the attendees on the Android Open Accessory Protocol



Amol Mujumdar, director engineering, Rediff.com, shared his views on 'Real-time Vulnerability Detection at Rediff'



Abhishek Dwivedi, head-capabilities, Vistaprint India, talked about application monitoring and log management



Sateesh Kavuri, senior architect, Yodlee, enlightened attendees on the MEAN stack



Ronen Baram, MySQL sales consultant, conducted a workshop on 'MySQL Performance Tuning'



Dushyant Min, CTO, Unotech Software Pvt Ltd, talked about moving fast with high performance Hack and PHP



Srinivas Acharya, engineering manager, HP Cloud, conducted a workshop on 'HP Helion OpenStack - a Technical Overview'



Pavithra Raman from Acquia Inc, conducted a workshop on 'Drupal in a Day'



Lux Rao, CTO, HP, talked about the transformation in the cloud, mapping the journey and the milestones



Rohit Rai, director, Diazle, gave a talk on 'Developing Commercially Viable Apps'



Ganesh Smartham, corporate trainer and consultant, shared his views on developing and marketing Web applications



Vikalp Sahni, CTO, Goibibo.com, talked on 'Goibibo: Technology Evolution, the Open Source Way'



Aveekshith Bhushan, solutions architect, MongoDB, talked about 'Building a Scalable and Modern Platform with MongoDB'



Pranav Kapoor from bookmyshow.com shared the journey of open source at BookMyShow



Jaison Justus from bookmyshow.com shared the journey of open source at BookMyShow



C N Krishnan, programme director, Anna University, shared his thoughts on FOSS adoption in the educational institutes in India

# An Introduction to Firefox OS

“We’re out to make a difference, not a profit,” says the project’s home page. The Firefox operating system has come into the market with a big bang. Read about how to install it on your computer system.



**F**irefox OS (project name: *Boot to Gecko*, also known as *B2G*) is a Linux kernel-based open source operating system for smartphones and tablet computers. It is being developed by Mozilla, the non-profit organisation best known for its Firefox Web browser. The Firefox OS is about pushing the boundaries of what is possible with the Web on the mobile, and enabling entirely new segments of users to come online with their first smartphones.

The Firefox OS is designed to provide a complete community-based alternative system for mobile devices, using open standards and approaches such as HTML 5 applications, JavaScript, a robust privilege model, open Web APIs to communicate directly with cell phone hardware and an application marketplace.

## Firefox OS terminology

There are a few terms that we need to get used to before we go further with the Firefox OS.

**B2G:** This is the short form for ‘Boot 2 Gecko’. This is the main engine on which the Firefox OS runs and is the engineering code name for the Firefox OS. You will often see this term used to refer to the Firefox OS, since it was used for a long time before the project had an official name.

**Firefox OS:** Firefox OS is basically Mozilla’s branding and support services applied on top of the B2G, to create a final release product.

**Gaia:** This is the user interface of the Firefox OS

platform. Gaia implements the lock screen, home screen and all the standard applications you expect on a modern smartphone. It is implemented entirely using HTML, CSS and JavaScript. Its only interfaces to the underlying operating system are through Web APIs that are implemented by the Gecko layer. Third party applications can be installed alongside the Gaia layer.

**Gecko:** This is the Firefox OS application runtime, i.e., this layer provides all the support for the triad of open standards—HTML, CSS and JavaScript. It makes sure these APIs work well on every operating system Gecko supports. This means that Gecko includes, among other things, a networking stack, graphics stack, layout engine, a JavaScript virtual machine and porting layers.

**Gonk:** Gonk is the lower level operating system of the Firefox OS platforms, consisting of a Linux kernel and user space hardware abstraction layer (HAL). The kernel and several of the user space libraries are common open source projects: Linux, libbus, Bluez, etc. Some of the parts of the HAL are shared with the AOSP (Android Open Source Project), such as GPS, camera and others. You could say that Gonk is a very simple Linux distribution. It is a porting target of Gecko, i.e., there’s a port of Gecko to Gonk, just like there’s a port of Gecko to OS-X, Windows and Android. Since the Firefox OS project has full control over Gonk, we can expose interfaces to Gecko that can’t be exposed on other operating systems.

**Jank:** This term is often used in the mobile apps space, and refers to the effect of slow or inefficient code operations in an app, which block the updating of the UI and cause it to lag or become unresponsive.

## Firefox OS build prerequisites

Here is a small guide on how to build the Boot 2 Gecko project's source code.

### Requirements for GNU/Linux

To build the B2G project source code on Linux, you need the following minimum requirements:

- A 64-bit GNU/Linux distribution
- At least 4GB RAM
- Free hard disk space of 30GB

This is more than the bare minimum requirements, but sometimes the build might fail just because of some missing resources. A typical error in this case is 'arm-linux-androideabi-g++. Internal error: Killed (program cc1plus)'.

You will also need the following tools installed in the system before you start the build process:

autoconf 2.13, bison, bzip2, ccache, curl, flex, gawk, git, gcc / g++ / g++-multilib, java sdk (jdk), make, OpenGL shared libraries, patch, X11 headers, 32-bit ncurses, 32-bit zlib

### Requirements for 64-bit installation

**Ubuntu/Linux Mint:** Run the following commands in a terminal:

```
sudo apt-get update
sudo apt-get install autoconf2.13 bison bzip2 ccache curl
flex gawk gcc g++ g++-multilib git ia32-libs lib32ncurses5-
dev lib32z1-dev libgl1-mesa-dev libx11-dev make zip
```

If you are building for the Flame reference phone or for Nexus 5, run the following command in a terminal:

```
Sudo apt-get install libxml2-utils
```

In addition to the emulator build issues discussed above, the compiler will default to *gcc-4.7*, which will fail to build due to an error along the following lines:

```
"KeyedVector.h:193:31: error: indexOfKey was not declared
in this scope, and no declarations were found by argument-
dependent lookup at the point of instantiation"
```

In a fresh Ubuntu 12.10 install, you'll get an error about unmet dependencies for *ia32-libs*. The following commands fix it:

```
sudo dpkg --add architecture i386
```

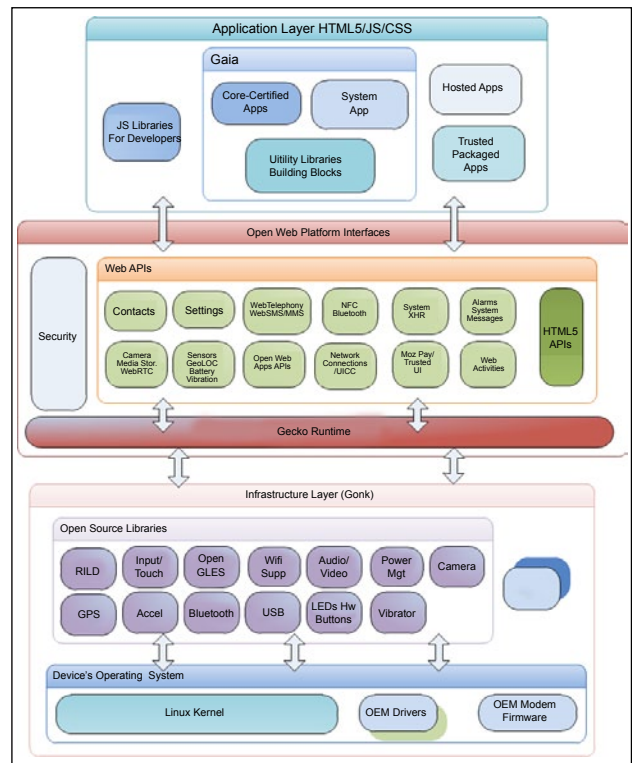


Figure 1: Architecture

```
sudo apt-get update
sudo apt-get install ia32-libs
```

**Fedora 19/20:** Those using Fedora 19/20 can type the following commands in a terminal:

```
sudo yum install autoconf213 bison bzip2 ccache curl flex
gawk gcc-c++ git glibc-devel glibc-static libstdc++-static
libx11-devel make mesa-libGL-devel ncurses-devel patch zlib-
devel ncurses-devel.i686 readline-devel.i686 zlib-devel.i686
libx11-devel.i686 mesa-libGL-devel.i686 glibc-devel.i686
libstdc++-i686 libXrandr.i686 zip perl-Digest-SHA wget
```

The project should build correctly using the default compiler shipped with Fedora but if you encounter compilation errors, you might need to install *GCC 4.6.x*:

Download the *tarball* and then install it to */opt* with the following command on Fedora 19/20:

```
curl url -O http://people.mozilla.org/~gsvelto/gcc-4.6.4-
fc19.tar.xz
sudo tar -x -a -C /opt -f gcc-4.6.4-fc19.tar.xz
```

**Arch Linux:** To install the required resources in Arch Linux, run the following command in the terminal:

```
sudo pacman -S --needed alsa-lib autoconf2.13 bison ccache
curl firefox flex gcc-multilib git gperf libnotify libxt
```



```
libx11 mesa multilib-devel wget wireless_tools yasm zip
lib32-mesa lib32-mesa-libgl lib32-ncurses lib32-readline
lib32-zlib
```

B2G can be only compiled with `gcc4.6.4`, and because Arch Linux always has bleeding edge software, you will need to install `gcc46-multilib` from AUR. Remember that you will have to edit the `PKGBUILD` and add `staticlibs` to the `options` array, or `gcc` will be unable to compile B2G and give you a `'cannot find -lgcc'` error when compiling. You will also need to add the following to your `.userconfig` file:

```
export CC=gcc-4.6.4
export CXX=g++-4.6.4
```

By default, Arch Linux uses Python3. You'll have to force it to use the old Python2. You can do that by linking the Python2 executable to Python, but this is discouraged since it is considered error-prone. This will also break Python3 if it is installed on your system. A better way is to use `virtualenv/virtualenvwrapper`:

```
sudo pacman -S python-virtualenvwrapper
source /usr/bin/virtualenvwrapper.sh
mkvirtualenv -p 'which python2' firefoxos
workon firefoxos
```

Android will complain that you need `make 3.81` or `make 3.82` instead of 4.0. You can download `make 2.81` from AUR. This will install the `make-3.81` binary on your path; you need to create a `symlink` named `make` by retaining the same location as mentioned in the `PATH` variable for the build to use the correct version:

```
mkdir -p ~/bin
ln -s 'which make-3.81' ~/bin/make
export PATH=~:/bin:$PATH
```

Android also needs the Java6 SDK, and Arch only has Java7. Unfortunately, the AUR build is broken, but you can still download the Java6 SDK and install it manually. You will then need to put it in your path.

```
cp ~/Downloads/jdk-6u45-linux-x64.bin/opt
su
cd /opt
chmod +x jdk-6u45-linux-x64.bin
./jdk-6u45-linux-x64.bin
exit
ln -s /opt/jdk1.6.0_45/bin/java ~/bin/java
```

**Gentoo Linux:** You need to install `ccache`, a tool for caching partial builds:

```
emerge -av ccache
```

Because `ccache` is known to frequently cause support issues, Gentoo encourages you to use it explicitly and sparingly.

To enable the required use of `ccache`, in the subsequent step of this guide in which the `.build.shscript` is called, Gentoo users should instead run the command with an explicitly extended path, i.e.:

```
PATH =/usr/lib64/ccache/bin:$PATH ./build.sh
```

## Install ADB

The build process needs to pull binary blobs from the Android installation on the phone before building B2G (unless you're building the emulator, of course). For this, you will need ADB, the Android Debug Bridge.



**Note:** Remember that when you start to use ADB, the phone's lock screen will need to be unlocked to view your phone's contents (at least in later versions of the Firefox OS). You'll probably want to disable the lock screen (we'll get to that later in the build instructions).

## Install Heimdall

Heimdall is a utility for flashing the Samsung Galaxy S2. It's used by the B2G flash utility to replace the contents of the phone with Firefox OS, as well as to flash updated versions of B2G and Gaia onto the device. You'll need it if you want to install Firefox OS on a Galaxy S2; it is not needed for any other device. For other devices, we build and use the fastboot utility instead.

There are two ways to install Heimdall:

- You can download the code from Github and build it yourself
- Use the package manager to install Heimdall
  - Run the following command in the terminal: `"sudo apt-get install libusb-1.0-0 libusb-1.0-0-dev"`

## Configuring ccache

The B2G build process uses `ccache`. The default cache size for `ccache` is 1GB, but the B2G build easily saturates this; so around 10GB is recommended. You can configure your cache by running the following command inside the terminal:

```
ccache -max-size 10GB
```

## Enabling remote debugging

Before you plug your phone back into your USB port, put it in USB developer mode. This allows you to debug and flash the phone. To enable developer mode, enable *Remote Debugging* in *Developer Settings*. Once the option is checked, remote debugging is enabled, and you are ready to go.

At this point, connect your phone to your computer via a

USB cable (if you created the *udev* rule earlier, this will trigger *udev* to detect the phone and create the device node with the right permissions). You can now check if you can list your device via the ADB devices command (remember that ADB can only see your phone when the lock screen is unlocked). If everything has worked well, you should see an output similar to the following (which is for a Geeksphone Keon):

```
$ adb devices
List of devices attached
Full_keon device
```

## Building Firefox OS

Depending on your Internet connection, the configuration takes a number of hours to download the files necessary to build Firefox OS (with a slow 150kbps connection, downloading gigabytes of Android repositories can take tens of hours). Waiting is not as much fun as actually doing something, so after you have read through this page and have kicked off the configure script, consider using the time to set up and try out the Firefox OS simulator. Begin familiarising yourself with 'Documentation for app developers' including 'Designing and Building an App'. You could also read up on information related to the upcoming steps.

## Cloning the B2G repository

The first step, before you can start your first build, is to clone the B2G repository. This will not fetch everything! Instead, it will fetch the B2G build system and set-up utilities. Most of the actual B2G code is in the main Mozilla *Mercurial* repository.

To clone this repository, use Git and run the following commands in the terminal:

```
git clone git://github.com/mozilla-b2g/B2G/git
cd B2G
```

Once you've retrieved the core B2G build system, you need to configure it for the device on which you plan to install it. To get a list of supported devices, you can use the *config.sh* utility; run the following command from within the B2G directory:

```
./config.sh
```

When you enter the above command, you will get something like the following in the terminal screen:

```
Usage: ./config.sh [-cdflnq] (device name)
Flags are passed through to |./repo sync|
```

Valid devices to configure are:

```
- galaxy-s2
```

```
- galaxy-s2
- nexus-4
- nexus-s
- nexus-s-4g
- flo "(Nexus 7 2013)"
- nexus-5
- flame
- otoro
- unagi
- inari
- keon
- peak
- leo
- hamachi
- helix
- wasabi
- fugu
- tara
- pandaboard
- emulator
- emulator -jb
- emulator -kk
- emulator -x86
- emulator -x86-jb
- emulator -x86-kk
```

Now, depending on the device you are working on, you can just build the emulator for it. For example, when you want to build the emulator for a Nexus device, you need to run the following command in the terminal:

```
./config nexus
```

## Updating your B2G tree

When the repository is updated to a newer version of B2G, you'll want to update your B2G tree. To do this, you can run the following commands in the terminal:

```
git fetch origin
git checkout origin/master
```

You can check that these worked correctly by running the following command:

```
git show HEAD
```

## Starting the build process

### Updating your code

If this isn't your very first time building B2G, you might want to pull the latest code before you start to build. To do that, you should update both the B2G tools and the dependencies, using the following two commands:

```
git pull
./repo sync -d
```

The `-d` option switches the various projects' `HEAD` back to their Android manifest revision (i.e., the default mainstream repository version). This is helpful if you made modifications to the repos (and have associated source changes) but want to revert back to the master version temporarily. It will leave your staged or working directory changes untouched.

## Building the B2G tree

To start the build process, run the following command in the terminal inside the B2G folder:

```
./build.sh
```

It will take around one hour or more to build the B2Ge source code locally. Once you are done building the source code, you can start the emulator by running the following command in the terminal in the B2G folder itself:

```
./run-emulator.sh
```


## Submitting bug reports on B2G, Firefox OS and Gaia

Once you've got a B2G/Firefox OS build running, you'll probably want to start to file bug reports against this specific

version, so that the Firefox OS community can improve things as effectively as possible. You'll need to file the bug on Bugzilla, under the 'Firefox OS' project, but you should include specifics about the version:

- To start with, give the major version number, e.g., 1.4.0.0-pre-release. This can be found on the device under *Settings > Device Information*
- You can provide more specific version identifiers by returning the current hashes from the Gaia and Gecko repositories. This can be done as follows:

```
#!/bin/bash
(cd gaia; echo "gaia $(git rev-parse HEAD)")
(cd gecko; echo "gecko $(git rev-parse HEAD)")
```

After making some changes in the B2G tree, you need to push them to the master tree, which is based on Git. **END** 

## References

- [1] [https://developer.mozilla.org/en-US/Firefox\\_OS/](https://developer.mozilla.org/en-US/Firefox_OS/)

## By Anup Kumar

The author is an active developer in Mozilla and GNOME. He regularly blogs at <https://anup07.wordpress.com/> and you can reach him at [allamsetty.anup@gmail.com](mailto:allamsetty.anup@gmail.com).

## ...Continued from page 25

Given a set of questions based on world events, and access to the World Wide Web (and its search engines), can you come up with a program that generates answers to the questions based on search query results? One possible way is to do a keyword-based search on the search query results and return the document passage that bears the most similarity. What are the other possible techniques you can use? For instance, consider the following two questions:


- What is the major food item exported by India?
- What is the only city in India which, apart from being the capital of a state, is also located on the banks of a river and is on the seashore?

How do these two questions vary in terms of difficulty in finding answers through an automated system? How would your solution deal with question (b)?

By the way, I wanted to point our readers to an interesting article I read recently in the *IEEE Spectrum* magazine by Prof Jordan (of machine learning, EM algorithm and Bayesian networks fame) titled 'Delusions of Big Data and Other Huge Engineering Efforts' which is available at:

<http://spectrum.ieee.org/robotics/artificial-intelligence/machinelearning-maestro-michael-jordan->

*on-the-delusions-of-big-data-and-other-huge-engineering-efforts*. He makes the important point that the inferences drawn from Big Data need to be validated for whether they are random patterns found by analysis (causality does not imply causation) or real root-causes explaining the data patterns. On being asked what research area he would target if he got a billion dollar research grant, he picked Natural Language Processing.

If you have any favourite programming questions/software topics that you would like to discuss on this forum, please send them to me, along with your solutions and feedback, at [sandyasm\\_AT\\_yahoo\\_DOT\\_com](mailto:sandyasm_AT_yahoo_DOT_com). Till we meet again next month, Happy New Year and happy programming! **END** 

## By: Sandya Mannarswamy

The author is an expert in systems software and is currently working with Hewlett Packard India Ltd. Her interests include compilers, multi-core and storage systems. If you are preparing for systems software interviews, you may find it useful to visit Sandya's LinkedIn group Computer Science Interview Training India at [http://www.linkedin.com/groups?home=HYPERLINK "http://www.linkedin.com/groups?home=&gid=2339182"&HYPERLINK "http://www.linkedin.com/groups?home=&gid=2339182"](http://www.linkedin.com/groups?home=HYPERLINK%20http://www.linkedin.com/groups?home=&gid=2339182%20HYPERLINK%20http://www.linkedin.com/groups?home=&gid=2339182)





In the above code, in Line No. 11, we are trying to change the value of the variable to which the pointer is 'pointing to', but this is not possible since the value is constant. When the above code is compiled and run, we get the output shown in Figure 1.

## Illustration 2

Now, let's use the same example given in Illustration 1 to show that the 'address' that the pointer contains is not a constant.

```
satya@satya:~$ gcc pointer_const.c
pointer_const.c: In function 'main':
pointer_const.c:12:2: error: assignment of read-only location '*ptr'
 *ptr = 30; //Error: Since, the pointer pointing to the value is constant.
  ^
satya@satya:~$
```

Figure 1: Output of the code snippet given in Illustration 1

```
1 #include <stdio.h>
2
3 int main()
4 {
5     //Definition of the variables
6     int a = 10;
7     int b = 20;
8
9     //Definition of pointer to constant
10    const int* ptr = &a; //Now, ptr is pointing to
    the value of the variable 'a'
11
12    ptr = &b; // Works: Since pointer is not constant
13
14    return 0;
15 }
```

From Illustrations 1 and 2, one can understand that the 'address' that the pointer contains can be changed but not the value to which the pointer is 'pointing to'. This can be clearly understood by the pictorial representations given in Figures 2, 3 and 4.

## Constant pointers

A 'constant pointer' is one that cannot change the address it contains. In other words, we can say that once a constant pointer points to a variable, it cannot point to any other variable.




Figure 2: Pictorial representation of 'pointer to constant'




Figure 3: Output of the code snippet given in Illustration 3




Figure 4: Pictorial representation of 'constant pointer'

Table 2: Pointer to constant concept

| Pointer to constant | Value change | Address change |
|---------------------|--------------|----------------|
| Const int*ptr;      | Not possible | Possible       |



**Note:** However, these pointers can change the value of the variable they 'point to' but cannot change the address they are 'holding'.

Table 3: Showing how to declare 'constant pointer'

| Syntax                                 | Example       |
|--|---------------|
| <type of pointer>*const <pointer name> | int*const ptr |

Table 4: Constant pointer concept

| Pointer to constant | Value change | Address change |
|---------------------|--------------|----------------|
| int*const ptr;      | Possible     | Not possible   |

Table 5: Summary

| Example              | Value constant | Pointer constant |
|----------------------|----------------|------------------|
| char*ptr             | No             | No               |
| const char*ptr       | Yes            | No               |
| char const*ptr       | Yes            | No               |
| char*const ptr       | No             | Yes              |
| const char*const ptr | Yes            | Yes              |

## Illustration 3

Let us consider the following code snippet to understand how 'constant pointer' works:

```
1 #include <stdio.h>
2
3 int main()
4 {
5     //Definition of the variable
6     int a = 10;
7     int b = 20;
8
9     //Definition of pointer to constant
10    const int* ptr = &a; //Now, ptr is pointing to the
    value of the variable 'a'
11
12    *ptr = 30; // Works, since the pointer pointing to
```

Table 6: Summary without asterisk

| Example              | Part Before Asterisk | Part After Asterisk | Comments   |
|----------------------|----------------------|---------------------|--|
| const char*ptr       | const                | ptr                 | Const is associated with data type, so value is constant               |
| char const*ptr       | char const           | ptr                 | Const is associated with data type, so value is constant               |
| char*const ptr       | char                 | const ptr           | Const is associated with pointer, so pointer is constant               |
| const char*const ptr | const char           | const ptr           | Const is associated with both data type & pointer so both are constant |

```

satya@satya: ~
satya@satya:~$ gcc pointer_const.c
pointer_const.c: In function 'main':
pointer_const.c:14:2: error: assignment of read-only variable 'ptr'
  ptr = &b; //Error :Now, ptr is pointing to value of the variable 'b'
  ^
satya@satya:~$

```

Figure 5: Output of the code snippet shown in Illustration 3

```

satya@satya: ~
Linux Programmer's Manual
strlen(3)
NAME
  strlen - calculate the length of a string
SYNOPSIS
  #include <string.h>
  size_t strlen(const char *s);
DESCRIPTION
  The strlen() function calculates the length of the string s, excluding
  the terminating null byte ('\0').
RETURN VALUE
  The strlen() function returns the number of bytes in the string s.

```

Figure 9: Shows the usage of pointer to constant in strlen() library function




Figure 6: Pictorial representation of constant pointer usage




Figure 7: Pictorial representation showing value contained in the variable can be changed through the constant pointer

```

satya@satya: ~
Linux Programmer's Manual
strcmp(3)
NAME
  strcmp, strncmp - compare two strings
SYNOPSIS
  #include <string.h>
  int strcmp(const char *s1, const char *s2);
  int strncmp(const char *s1, const char *s2, size_t n);
DESCRIPTION
  The strcmp() function compares the two strings s1 and s2. It returns an integer less than, equal to, or greater than zero if s1 is found, respectively, to be less than, to match, or be greater than s2.
  The strncmp() function is similar, except it compares the only first (at most) n bytes of s1 and s2.
RETURN VALUE
  The strcmp() and strncmp() functions return an integer less than, equal to, or greater than zero if s1 (or the first n bytes thereof) is found, respectively, to be less than, to match, or be greater than s2.

```

Figure 10: Shows the usage of pointer to constant in strcmp() library function




Figure 8: Pictorial representation showing constant pointer value can not be changed

```

the value is not constant
13
14     ptr = &b; //
Error:Now, ptr is pointing to
the value of the variable 'b'
15
16     return 0;
17
18 }

```

From the above example (Illustration 3), it is clear that in Line No 14 we tried to change the address of the pointer *ptr* to some other variable, but it is not possible. The output of the code snippet shown in Illustration 3 is given in Figure 5. Similarly, one can observe that in Line No 12, we are trying to change the value of the variable it is 'pointing to', which is possible.

This can be clearly understood by the pictorial representations given in Figures 6, 7 and 8.

### Something to think about

Can we have both pointer to constant and constant pointer in a single statement?

## Usage

We can find 'n' number of uses of these concepts in C as well as in the embedded C programming world. One

such simple use of 'pointer to constant' is to find the string length of the given string without any attempt to modify the original string as shown in Example 1 (Figure 9). Example 2 gives an idea of using 'pointer to constant' in the *strcmp()* function (Figure 10).

## A trick

There is a small trick to understand the difference between 'pointer to constant' and 'constant pointers' which is shown in Table 6.



**Note:** This trick is for all those new to the C programming world, who are confused with constant and pointers.

From the summary shown in Table 5, separate the part before asterisk(\*) and the part after the asterisk(\*) as given in Table 6, to clearly understand whether 'data' is constant or 'pointer' is constant. **END**

## By: Satyanarayana Sampangi

The author is a member - Embedded software at Emertxe Information Technologies (<http://www.emertxe.com>). His area of interest lies in embedded C programming combined with data structures and microcontrollers. He likes to experiment with C programming and open source tools in his spare time to explore new horizons. He can be reached at [satya@emertxe.com](mailto:satya@emertxe.com)



# GUARD YOUR WEB APPLICATIONS AGAINST THE **TOP 10 VULNERABILITIES**



A flawed Web application is a security risk. OWASP's 'Top 10 Vulnerabilities' is a powerful document that raises awareness about the dangers out there and pinpoints possible flaws in a Web application. It also suggests ways and means to counter these susceptibilities. OWASP enjoys the support and backing of security experts from around the world.

Allow me to begin this article by asking a simple question. Can you say, with absolute surety, that your house is theft-proof? If your answer is, "Yes," then you're living under a false sense of security. However, if your answer is "No," then we have something to talk about. You may have state-of-the-art security systems installed in your home but none will be a good enough match against determined burglars. The security system will definitely make their task difficult, but it won't be able to stop them. All that would be required is for them to find the Achilles heel of your security system. You might wonder that if no security system is good then why even bother installing one? Why not leave your house wide open for the burglar? While security systems might not block out the burglars entirely, they will enable you to protect the vulnerable spots of your house and give you enough time to detect the burglars and take action.

If you replace the word 'house' with 'Web application' in the above scenario, the argument is still valid. The only difference is that the number of vulnerable spots in a Web application is much more than in a house. Also, given the fact that a Web application will most likely be exposed to the world at large, the number of threat agents and attack vectors increases exponentially. Therefore, in order to develop a secure Web application, developers will have to think of each of the possible ways in which their app can be compromised. Considering that there are more than 500,000 ways in which this can be done, by the time the application is ready to hit the market after being tested, it might already be out-of-date. How do developers then ensure that they develop a secure application without any significant impact on its time to market and its usability? The answer to this lies in the Open Web Application Security Project (OWASP).

## The Open Web Application Security Project (OWASP)

The Open Web Application Security Project (OWASP) (<http://owasp.org>) is a not for profit community dedicated to help developers design and develop secure Web applications. It enables them to prioritise their efforts by publishing various standards, guides, cheat sheets, etc, in order to secure applications. Started in 2001 by Mark Curphey, it has grown to include several organisations, educational institutions and volunteers working towards building a secure Web.

Among its various projects, every three years, OWASP publishes a list of the top 10 vulnerabilities that plague Web applications. The list is published after extensive research, which includes the survey of the top four consulting companies in the world and three software vendors. In total, the project goes through a database of approximately 500,000 vulnerabilities and shortlists the top 10. The list was last published in 2013.

In addition to publishing this list, OWASP also includes the means and methods to counter them. The following section briefly describes each of the 10 vulnerabilities and their countermeasures (for more details, please visit <http://goo.gl/p6rAzr>).

## OWASP's top 10 vulnerabilities, 2013

### 1. Injection

A Web application is vulnerable to injection when it accepts commands or queries in input fields, meant for obtaining information from the user. A classic example is SQL Injection, wherein the attacker injects SQL queries in an input field in order to bypass the authentication mechanism.

*How to fix it:* Use safe or parametrised APIs, escape special characters, provide white list input validation (e.g., if a field is meant to accept numeric values, the application should not permit the user to enter letters of the alphabet).

### 2. Broken authentication and session management

This flaw can occur when authentication mechanisms are not implemented properly (e.g., sending or storing credentials in plain text; when password recovery allows passwords to be changed without proper authentication and verification of the user, etc) or when sessions are poorly managed (e.g., if the session time-out is not defined; if session IDs are exposed in the URL, etc).

*How to fix it:* Use standards like ISO 27001:2013 or OWASP's Application Security Verification Standard (ASVS) when defining authentication and session management requirements for the Web application. Ensure that the Cross Site Scripting (XSS - explained later) flaw is taken care of.

### 3. Cross Site Scripting (XSS)

A Cross Site Scripting vulnerability allows an attacker to insert malicious script/code in a Web page that directs the user's browser to perform a malicious

action. For example, an attacker can leave a simple JavaScript code in a website's comments section, which redirects anyone visiting that website to a fake login page.

*How to fix it:* Avoid special characters, use white list input validation, use auto-sanitisation libraries for rich content, and employ the Content Security Policy to protect the entire website from XSS.

### 4. Insecure direct object references

Web applications often expose actual object names to users while generating Web pages (e.g., `<URL>/acctID='1234'`). However, if they fail to verify a user's access privilege to that particular object, users could manipulate the object value and access information that they are not authorised to. For example, in the above URL, a user may put acctID as '4567' and be able to access the information of that account even though he's not privileged to do so.

*How to fix it:* Check users' access privileges to each object reference before granting access to it; use session-specific or user-specific mapping of objects to avoid direct references.

### 5. Security misconfiguration

Most software like Web servers, database servers, programming platforms, etc, ship with disabled security controls. Often, Web developers either forget to configure

Want Latest News?  
Check out  
**EFYTimes.com**



Want Industry Trends?  
Check out  
**Electronicsb2b.com**



An  
**EFYGROUP**  
PROPERTY



## 2 Simple Ways

# To Stay Up-to-date

## With India's

# Electronics Industry

For any queries, please contact our team at [efyenq@efy.in](mailto:efyenq@efy.in) OR +91-11-26810601

these controls or configure them loosely. This enables the attacker to break into the system very easily.

*How to fix it:* Harden the system, i.e., disable unnecessary functionality in the underlying applications, patch and update the systems regularly, and keep track of the libraries used in the Web application (update, if the new version is available or discard if it's no longer supported). Regularly run vulnerability scans and penetration tests to uncover flaws.

#### 6. Exposure to sensitive data

One of the most common flaws in Web applications is their failure to protect sensitive data adequately. They either store or transmit data in plain text, use weak encryption and hashing algorithms, or the keys aren't properly managed. Exploitation of all or any one of these can have a massive negative impact on a business' reputation.

*How to fix it:* Use FIPS-certified data encryption algorithms, collect and store only what's necessary, delete the data once it has served its purpose and disable auto-complete.

#### 7. Missing function-level access control

This flaw occurs when the access to application functions isn't validated. For example, a Web application may hide the link to the *Administration* console by not presenting it to a normal user. However, it does not implement any function level access control to check who is accessing administration functions. As a result, normal users can access administration functions, provided they know which ones to use.

*How to fix it:* Implement access control on all functions handling sensitive data or functionality. At the least, assign roles to the users and validate their roles before granting access to sensitive functions.

#### 8. Cross-Site Request Forgery (CSRF)

CSRF occurs when an attacker can craft a Web page that requests the user's browser to send session information of some other website (e.g., *sample.com*). The browser senses this as a valid request and sends the requested information to the attacker's server. Thus the attacker can forge the user's session on *sample.com* with the stolen information.

*How to fix it:* Hide CSRF tokens in the body of the Web page, use CAPTCHA as a means for users to confirm their actions, etc.

#### 9. Using components with known vulnerabilities

Most Web applications use third party libraries to implement certain functionality. While this saves on development efforts and time, it may become a developer's nightmare if such libraries are not managed properly. For example, consider a Web application using a certain reporting library, which has reached its end-of-support phase. If the developers fail to keep track of this library, their application will be vulnerable to future vulnerabilities affecting the library.

*How to fix it:* Keep track of the versions, patches and upgrades of the libraries in use; discard libraries that are no longer in development or have reached the end-of-support

phase. Create a policy outlining the process for safe usage of such libraries.

#### 10. Invalid redirects and forwards

Redirects and forwards are a common occurrence in a Web application and it's important to check whether or not they lead to a valid or correct page. If attackers can modify the parameters of the redirection URL and change the landing page, they can bypass the website's security mechanisms or lead users to disclose sensitive information.


*How to fix it:* Avoid using redirects and forwards. If they can't be avoided, ensure that no user parameters are involved in generating the landing page. If user parameters are to be used, ensure that they are properly validated.

### Additional measures to secure Web applications

OWASP's ranking of the top ten vulnerabilities enables developers to secure their Web applications against the most prevalent vulnerabilities but this complex task doesn't end here. Once these are taken care of, Web developers should look at the environments in which their application will be used, who are its target users, etc. For example, if the application is being developed for military purposes, it'll need to have strong access control mechanisms. If it is to be used in the health care industry, then encryption and the protection of data from leakage takes priority.

The IT environment differs from one organisation to another and each environment has its set of vulnerabilities. Web applications developers must take this factor into consideration and secure their applications with respect to the environment in which the Web application will operate.

Protecting applications from hobby hackers is a thing of the past. Attackers now have advanced skill sets and use sophisticated techniques to crack websites. Day in and day out there's news of some website being hacked and sensitive data being leaked (e.g., Dropbox, Apple iCloud, RSA, Snapchat, etc). Cracking into websites has become a booming business in the online underworld.


While attackers only have to find a new attack vector, Web developers are tasked with protecting their applications against known attacks as well as those that are yet unknown. The OWASP community has done a lot to ease the burden of Web developers by publishing the list of the top 10 vulnerabilities. Based on this, developers can prioritise their development efforts and develop secure Web applications without compromising a project's deadline or the overall user experience. 

**By: Uday Mittal**

The author is an information security consultant, an open source enthusiast and likes to experiment with new technologies. He can be reached at [mailme@udaymittal.com](mailto:mailme@udaymittal.com)



**Goal #1**  
**Promote**  
**'Make**  
**in India'**



**Goal #2**  
**Showcase**  
**'Latest in**  
**Electronics'**

5th EDITION

**Electronics**  
**ForYou expo** 2015  
Innovate. Design. Manufacture. Source. **INDIA**

February 26 to 28, 2015  
Pragati Maidan, New Delhi

Organised by the Electronics For You Group, EFY Expo India is a "Made in India" show, which has only 2 goals: (1) to promote 'Make in India' and (2) to showcase 'Latest in Electronics'. If you support either of these two goals, be there.

**HOW TO 'BE THERE'?**

• To get a **Business Invite**, register at [www.efyexpo.com](http://www.efyexpo.com) OR contact us at [efyexpo@efy.in](mailto:efyexpo@efy.in) OR at +91-11-4059 6605

• To know how **you can increase your sales** by exhibiting at EFY Expo India, contact us as [growmybiz@efy.in](mailto:growmybiz@efy.in) OR at +91-88000 94213.

**HURRY! TO AVAIL PREMIUM LOCATION CONTACT: Praveen 088000 94207 OR Arun 088000 94213**

**POWERED BY**

Associate Partners

**TC Components**  
Quality is Growth

**MetroQ**

**MECHONICS**

**JUKI**  
JUKI INDIA PVT. LTD

**Tektronix**

**ASYS GROUP**  
Digital Circuits Pvt. Ltd.

**ASYS GROUP**

**emst**  
Marketing Pvt. Ltd.

**SGS Teknics**

# Enhance Java Applications with FOSS APIs

Java application developers have access to many free and open source APIs that ease their workload. This article, which presumes the reader has some knowledge of Netbeans, Core Java and Java Swing, demonstrates the implementation of user input validation, reporting and SMS integration.



Imagine a world where car manufacturers had to invent the wheel for every car they manufactured, or if Jonas Salk, the inventor of the polio vaccine, had patented his invention or if Dennis Ritchie, the creator of the C programming language, had not created it. What would it be like then? Very different from what it is today, wouldn't it? Thanks to the open source community, we aren't living in such a world.

The Free and Open Source Software (FOSS) community plays an important role in the world of software development. Over the past many years, its efforts and expertise have helped software developers to avoid re-inventing the wheel. Most of the software being developed today uses certain functionalities that have been implemented in some other software. In order to save development time, software vendors break such functionalities into modules and offer them as APIs. They may or may not charge other developers for using these APIs. Those who don't, often make them open source and invite the community to develop them further.

In this article, we explore three such open source APIs for developing Java applications—the SimpleValidation API, SMSLib API and DynamicReports API. These APIs enable developers to implement user input validation controls, reporting and SMS integration without much overhead.

The following sections cover each of the three APIs in detail. First, they give a brief introduction to the API, followed by its download and implementation details. A sample snippet is included wherever feasible. Each section includes a link to the working example referred to.

## The SimpleValidation API

When I was in college, our software engineering professor asked us to remember this simple principle while developing software: 'Always assume that your software will be used by monkeys.' He didn't mean it as an offence, but the message he wanted to convey was that user input validation is mission critical in any software.

An application with relaxed user input validation not only scores badly on the quality front, but is also a security hazard. Therefore, developers tend to spend more time on coding controls for erroneous inputs than on coding actual functionality, irrespective of the programming language or platform used.

Thankfully, for Java application developers, there's good news in the form of the SimpleValidation API, which was developed by Tim Boudreau. It is a Java library developed to ease the coding of user input validation controls. It reduces the time required to code validation controls, thereby allowing developers to enhance the application further, without compromising on project deadlines. Through this library, developers can simply plug the required validation control in most of the swing components. It comes pre-built with some of the common validation controls like the non-empty string, numeric-only, URL, email, length of the input, etc. A more elaborate introduction to and the documentation for the API are both available at <http://goo.gl/MqeQii>.

## Download and build

The most recent source code of the API is available at [git://kenai.com/simplevalidation~git](https://github.com/kenai/simplevalidation). In order to use this API, you'll need to

check out the source files from the Git repository and build them with NetBeans. You can use the following steps to do so:

- Download and install Git for Windows (<http://git-scm.com/download/win>)
- Launch the Git GUI and select *Clone existing repository*
- In the source location, enter the above mentioned URL for the Git repository and specify the target directory on your machine and click *Clone*.
- Once the repository is cloned, open it in NetBeans and build it. The build process will generate *.jar* files, which can then be included in any project.

However, if you think that's too much of work, I have uploaded a compiled version of the API at <http://goo.gl/Olc5MN>. Download the zip file, extract it and follow the next section.

## Implementation

If you're using the GIT repository, copy the *simplevalidation-standalone-<version>.Jar* file from *standalone\target* folder and *nbstubs-<version>.jar* from *\nbstubs\target* folder to the project folder. However, if you have downloaded the compiled APIs from the above link, copy the two JAR files into the project folder.

Next, import the API in Netbeans. To do this, right click the project, select *Properties > Libraries > Add Jar/Folder*, specify the location of JAR files and press OK.

This API allows you to define the severity levels (FATAL, WARNING and INFO) of validation controls. The FATAL severity level is the strictest and requires the user to correct the input before submitting the form. The WARNING and INFO severity levels are a bit relaxed and can be used for validation of optional fields. It also decorates the Swing components (see Figure 1) according to the severity level set. The default is FATAL for all Swing components.

## Sample code

A sample code snippet to enable validation in a login form with a username field, a password field and a login button is given below:

```
private ValidationPanel validationPanel;
private ValidationGroup validationGroup;

validationPanel = new ValidationPanel();
validationGroup = validationPanel.getValidationGroup();
validationPanel.validate();

/*Username field validation (Non-Empty String)*/

validationGroup.add(textField1,StringValidators.REQUIRE_NON_
EMPTY_STRING);

/*Password field validation (Minimum Length 8, Non-
Empty String, Disallowed
Characters)*/

validationGroup.add(passwordField1,StringValidators.REQUIRE_NON_ }
```

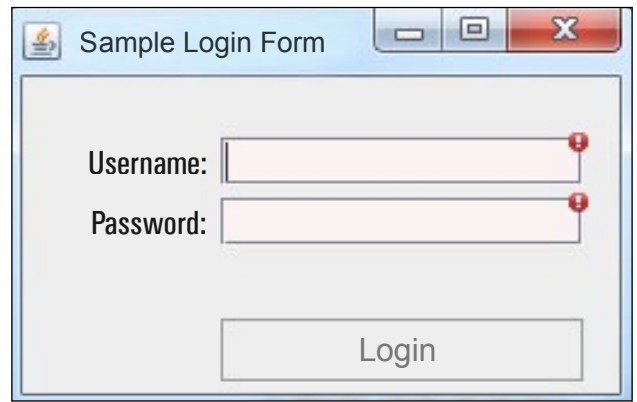


Figure 1: Login form

```
EMPTY_STRING,StringValidators.minLength(8),StringValidators.
disallowChars(new char[]{'\'','!'}));
validationGroup.addUI(new ValidationUpdate(button1));
```

```
/*Class to change the state of Login button depending on the
validity of input values and assigned severity*/
```

```
public class ValidationUpdate implements ValidationUI {
```

```
    JButton button,button2;
```

```
    public ValidationUpdate(JButton Button)
    {
        button=Button;
        button2=new JButton();
    }
```

```
    public ValidationUpdate(JButton Button, JButton Button2)
    {
        button=Button;
        button2=Button2;
    }
```

```
    public void clearProblem()
    {
        button.setEnabled(true);
        button2.setEnabled(true);
    }

    public void showProblem(Problem problem)
    {
        if(problem.isFatal())
        {
            button.setEnabled(false); button2.setEnabled(false);
        }
        else
        {
            button.setEnabled(true); button2.setEnabled(true);
        }
    }
}
```



You can download the code for the above example at <http://goo.gl/sIVtDb>. Extract the zip file, open the project in NetBeans, add the API JAR files to the project as mentioned above and run the project.

## The SMSLib API

A common requirement from clients is to integrate SMS functionality in their application. They may need it for sending order notifications, automated greetings, marketing messages, etc. Often this can be done using the bulk SMS gateways via the Internet. However, sometimes, using such gateways might not be a desired option. Clients may want to use their mobile or a GSM dongle to send such SMS. This method has two benefits—it is more personalised and it is not affected by 'Do Not Disturb' (DND) filters.

The SMSLib API, developed by Thanasis Delenikas, is an open source library, available for Java and Microsoft .NET platforms. Through this API, applications can send and receive messages using almost any GSM modem. It's simple to implement and can interface with most of the USB GSM dongles available in the market.

## Download

You have two options to use SMSLib in your projects. The first one is to download the source code of SMSLib v 3.5.x from <https://github.com/smslib/smslib-v3> and build it locally. If you choose this, follow the instructions available at <http://smslib.org/doc/installation/>. The second option is to download the compiled JAR file from <http://goo.gl/wx5oIZ> and import it into your NetBeans project using the steps mentioned previously.

## Implementation

Since SMSLib uses the Java Communication API, you'll have to download and install the Java Comm API before using it. The following instructions will help you in downloading and installing the Java Comm API:

- Download the JavaComm v2 (for Windows 32-bit systems) from <http://smslib.org/download/> and extract the zip file.
- Copy *Comm.jar* to:
  - %JAVA\_HOME%\lib (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\lib)
  - %JAVA\_HOME%\jre\lib\ext (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\jre\lib\ext)
- Copy *win32com.dll* to:
  - %JAVA\_HOME%\bin (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\bin)
  - %JAVA\_HOME%\jre\bin (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\jre\bin)
  - %windir%\System32 (e.g., C:\Windows\System32)
- Copy *javax.comm.properties* to:
  - %JAVA\_HOME%\lib (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\lib)

- %JAVA\_HOME%\jre\lib (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\jre\lib)



**Note:** The 'Program Files' (x86) folder exists only on 64-bit Windows machines.

Unfortunately, the JavaComm API is not available for 64-bit JDK; so in case you're using the 64-bit JDK, you'll have to use the RXTXComm API, which is an alternative to the JavaComm API and has a 64-bit version. The following instructions will help you in downloading and installing the RXTXComm API:

- Download the RXTXComm 64-bit version from <http://goo.gl/wx5oIZ> and extract the zip file.
- Copy *RXTXcomm.jar* to:
  - %JAVA\_HOME%\lib (e.g., C:\Program Files\Java\jdk1.8.0\_11\lib)
  - %JAVA\_HOME%\jre\lib\ext (e.g., C:\Program Files\Java\jdk1.8.0\_11\jre\lib\ext)
- Copy *rxtxSerial.dll* and *rxtxParallel.dll* to:
  - %JAVA\_HOME%\bin (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\bin)
  - %JAVA\_HOME%\jre\bin (e.g., C:\Program Files (x86)\Java\jdk1.8.0\_11\jre\bin)
  - %windir%\System32 (e.g., C:\Windows\System32).

## Sample code

In the sample code snippets shown below:

- Replace the 'Gateway Name' with 'GSM modem gateway' (usually it's the manufacturer's name).
  - Replace the COM port with the COM port on which your device is registered. You can get the list of registered COM ports from *Computer > Properties > Device Manager > Ports (COM & LPT)*. Try each port, if multiple ports are listed there.
  - Change the baud rate according to your device.
- Sample code to send an SMS is:

```
SerialModemGateway smsGateway = new
SerialModemGateway("Huawei", "COM8", 9600, "", "");
smsGateway.setInbound(true);
smsGateway.setOutbound(true);

try {
    Service.getInstance().addGateway(smsGateway);
    Service.getInstance().startService();

    /*Insert the recipient phone number here*/

    OutboundMessage msg = new OutboundMessage("+91<Insert Phone
    Number here>", "Test Message");

    Service.getInstance().
sendMessage(msg);
    Service.getInstance().stopService();
}
```

```

    } catch (GatewayException ex) {
    } catch (SMSLibException ex) {
    } catch (IOException ex) {
    } catch (InterruptedException ex) {
    }

```

Sample code to receive an SMS is:

```

SerialModemGateway gateway = new SerialModemGateway("Huawei",
"COM8", 115200, "", "");
gateway.setProtocol(Protocols.PDU);
gateway.setInbound(true);
gateway.setOutbound(true);
Service.getInstance().setInboundMessageNotification
(inboundNotification);
Service.getInstance().addGateway(gateway);
Service.getInstance().startService();
msgList = new ArrayList<InboundMessage>();
Service.getInstance().readMessages(msgList,
MessageClasses.ALL);

for (InboundMessage msg : msgList)
    System.out.println(msg);

System.out.println("Waiting for SMS. Press any key to stop.");
System.in.read();

/*Class to display the notification and text of incoming
message*/

public class InboundNotification implements
IInboundMessageNotification
{
    public void process(AGateway gateway, MessageTypes msgType,
InboundMessage msg)
    {
        if (msgType == MessageTypes.
INBOUND)
        {
            System.out.println("Incoming message from: " + gateway.
getGatewayId());
            System.out.println("Sender: " + msg.
getOriginator());
        }

        System.out.println(msg);
    }
}

```

The entire code for the above mentioned snippets is available at <http://goo.gl/uQDsCO> as a NetBeans project.

## The DynamicReports API

I came across the DynamicReports API when I was

developing my first business application. I was stuck for a while on the reporting module, thinking that I'd need to do it from scratch. That wasn't a motivating thought at all, so I searched the Internet for an alternative. I came across several proprietary reporting APIs before I stumbled upon DynamicReports, and since then I haven't looked anywhere else to create reports.


The DynamicReports API, developed by Ricardo Mariaca, is based on Jasper Reports. It reduces the effort required to create reports programmatically by abstracting much of the work required when using the Jasper Reports API. It comes pre-built with several designs and templates, thus eliminating the need for a visual designer. It also allows you to include various charts (bar charts, pie charts, metre charts, etc) in reports. It also allows you to export the reports in various formats (.pdf, .xls, .doc, etc). Thus, apart from some basic formatting and source data, it takes care of pretty much everything. You can visit <http://www.dynamicreports.org> to read more about it. It is well documented and the community provides good support.

## Download and implementation

You can download the API from <http://goo.gl/9xwY34>. At the time of writing, the version is 3.2.1. Once the zip file is downloaded, extract it and copy the contents of the *dist* folder to your project. Next, import the JAR files in your NetBeans project using the steps mentioned previously.

DynamicReports hosts a variety of examples at <http://www.dynamicreports.org/examples> to get you started. In case you need help, you can visit the support forum at <http://www.dynamicreports.org/forum/> and post your queries there or contact me at the email ID given at the end of the article.

The above mentioned APIs can be used in a variety of applications, either alone or in combination. For example, almost any application that requires users to input some value can benefit from the SimpleValidation API. DynamicReports can enhance any report-intensive business application (MIS, ERP, school results management, etc). SMSLib is very convenient for applications that use SMS as a means to accept user input (like accepting orders via SMS), and applications designed for machine-to-machine communication (for example, applications to monitor radio taxis, meter reading applications, etc).

Open source APIs and their developers are a boon to the developer community. They not only help in reducing development time and efforts but also bring down the overall cost of the project. **END** 

**By: Uday Mittal**

The author is an information security consultant, an open source enthusiast and a technical poet. He can be reached at [mailme@udaymittal.com](mailto:mailme@udaymittal.com), <http://genemesis.wordpress.com>.

# What is the Haskell I/O?

This article is a must read for anyone interested in getting a good insight into the input/output (IO) functionality of Haskell.



Input/output (IO) can cause side-effects and hence is implemented as a Monad. The IO Monad takes some input, does some computation and returns a value. The IO action is performed inside a main function. Consider a simple 'Hello world' example:

```
main = putStrLn "Hello, World!"
```

Executing the above code in GHCi produces the following output:

```
$ ghci hello.hs
GHCi, version 7.6.3: http://www.haskell.org/ghc/? for help
Loading package ghc-prim ... linking ... done.
Loading package integer-gmp ... linking ... done.
Loading package base ... linking ... done.
[1 of 1] Compiling Main ( foo.hs, interpreted )
Ok, modules loaded: Main.
ghci> main
Hello, World!
```

The type signatures of *main* and *putStrLn* are:

```
main :: IO ()
```

```
putStrLn :: String -> IO ()
```

*putStrLn* takes a string as input and prints the string to output. It doesn't return anything, and hence the return type is the empty tuple ().

The *getLine* function performs an IO to return a string.

```
ghci> :t getLine
getLine :: IO String
```

```
ghci> name <- getLine
Foo
```

```
ghci> name
"Foo"
```

The '<- ' extracts the result of the IO string action, unwraps it to obtain the string value, and 'name' gets the value. So, the type of 'name' is:

```
ghci> :t name
name :: String
```

The *do* syntax is useful to chain IO together. For example:



```
main = do
  putStrLn "Enter your name:"
  name <- getLine
  putStrLn ( "Hello " ++ name)
```

Executing the code in GHCi gives the following results:

```
ghci> main
Enter your name :
Shakthi
Hello Shakthi
```

The *putStr* function is similar to the *putStrLn* function, except that it doesn't emit the new line after printing the output string. Its type signature and an example are shown below:

```
ghci> :t putStr
putStr :: String -> IO ()

ghci> putStr "Alpha "
Alpha ghci>
```

The *putChar* function takes a single character as input, and prints the same to the output. For example:

```
ghc> :t putChar
putChar :: Char -> IO ()

ghci> putChar 's'
s
```

The *getChar* function is similar to the *getLine* function, except that it takes a *Char* as input. Its type signature and usage are illustrated below:

```
ghci> :t getChar
getChar :: IO Char

ghci> a <- getChar
d

ghci> a
'd'

ghci> :t a
a :: Char
```

The *print* function type signature is as follows:

```
ghci> :t print
print :: Show a => a -> IO ()
```

It is a parameterised function, which can take an input of any type that is an instance of the *Show* type class and prints

that to the output. Some examples are given below:

```
ghci> print 1
1

ghci> print 'c'
'c'

ghci> print "Hello"
"Hello"

ghci> print True
True
```

The *getContents* function reads the input until the end-of-file (EOF) and returns a string. Its type signature is shown below:

```
ghci> :t getContents
getContents :: IO String
```

An example of code is demonstrated below. It only outputs lines whose length is less than five characters:

```
main = do
  putStrLn "Enter text:"
  text <- getContents
  putStr . unlines . filter (\line -> length line < 5) $
    lines text
```

Testing the above example gives the following output:

```
ghci> main
Enter text:
a
a

it
it

the
the

four
four

empty

twelve

haskell
```

```
0
0
```

You can break out of this execution by pressing *Ctrl-C* at the GHCi prompt.

The *openFile*, *hGetContents*, *hClose* functions can be used to obtain a handle for a file, to retrieve the file contents, and to close the handle respectively. This is similar to file handling in C. Their type signatures are shown below:

```
ghci> :m System.IO

ghci> :t openFile

openFile :: FilePath -> IOMode -> IO Handle

ghci> :t hGetContents
hGetContents :: Handle -> IO String

ghci> :t hClose
hClose :: Handle -> IO ()
```

The different IO modes are *ReadMode*, *WriteMode*, *AppendMode* and *Read-WriteMode*. They are defined as follows:

```
-- | See 'System.IO.openFile'
data IOMode = ReadMode | WriteMode | AppendMode |
ReadWriteMode
deriving (Eq, Ord, Ix, Enum, Read, Show)
```

An example code is illustrated below:

```
import System.IO

main = do
  f <- openFile "/etc/resolv.conf" ReadMode
  text <- hGetContents f
  putStr text
  hClose f
```

Executing the code in GHCi produces the following output:

```
ghci> main

# Generated by NetworkManager
nameserver 192.168 .1 .1
```

A temporary file can be created using the *openTempFile* function. It takes as input a directory location, and a pattern string for the file name. Its type signature is as follows:

```
ghci> :t openTempFile
openTempFile :: FilePath -> String -> IO (FilePath, Handle)
```

An example is shown below:

```
import System.IO
import System.Directory (removeFile)
main = do
  (f, handle) <- openTempFile "/tmp" "abc"
  putStrLn f
  removeFile f
  hClose handle
```

You must ensure that you remove the file after using it. An example is given below:

```
ghci> main
/tmp/abc2731
```

The operations on opening a file to get a handle, getting the contents and closing the handle can be abstracted to a higher level. The *readFile* and *writeFile* functions can be used for this purpose. Their type signatures are as follows:

```
ghci> :t readFile
readFile :: FilePath -> IO String

ghci> :t writeFile
writeFile :: FilePath -> String -> IO ()
```

The */etc/resolv.conf* file is read and written to */tmp/resolv.conf* in the following example:

```
main = do
  text <- readFile "/etc/resolv.conf"
  writeFile "/tmp/resolv.conf" text
```

You can also append to a file using the *appendFile* function:

```
ghci> :t appendFile
appendFile :: FilePath -> String -> IO ()
```

An example is shown below:

```
main = do
  appendFile "/tmp/log.txt" "1"
  appendFile "/tmp/log.txt" "2"
  appendFile "/tmp/log.txt" "3"
```

The content of */tmp/log.txt* is '123'.

The actual definitions of *readFile*, *writeFile* and *appendFile* are in the *System.IO* module in the Haskell base package:

```
readFile :: FilePath -> IO String
readFile name = openFile name ReadMode >>= hGetContents

writeFile :: FilePath -> String -> IO()
writeFile f txt = withFile f WriteMode (\hdl-> hPutStr hdl txt)
```

```
appendFile :: FilePath -> String -> IO ()
appendFile f txt = withFile f AppendMode (\hdl ->
hPutStr hdl txt)
```

The *System.Environment* module has useful functions to read command line arguments. The *getArgs* function returns an array of arguments passed to the program. The *getProgName* provides the name of the program being executed. Their type signatures are shown below:

```
ghci> :m System.Environment
```

```
ghci> :t getArgs
getArgs :: IO [String]
```

```
ghci> :t getProgName
getProgName :: IO String
```

Here is an example:

```
import System.Environment
main = do
  args <- getArgs
  program <- getProgName
  putStrLn ( "Program : " ++ program)
  putStrLn "The arguments passed are: "
  mapM putStrLn args
```

Executing the above listed code produces the following output:

```
$ ghc--make args.hs

[1 of 1] Compiling Main ( args.hs, args.o )
Linking args ...

$ ./args 1 2 3 4 5

Program : foo
The arguments passed are:
1
2
3
4
5
```

The *mapM* function is the map function that works for Monad. Its type signature is:

```
ghci> :t mapM
mapM :: Monad m => (a -> m b) -> [a] -> m [b]
```

The *System.Directory* module has functions to operate on files and directories. A few examples are shown below:

```
ghci> :t createDirectory
createDirectory :: FilePath -> IO ()
```

```
ghci> createDirectory "/tmp/foo"
ghci>
```

If you try to create a directory that already exists, it will return an exception:

```
ghci> createDirectory "/tmp/bar"
*** Exception : /tmp/bar : createDirectory : already
exists (File exists)
```

You can use the *createDirectoryIfMissing* function, and pass a Boolean option to indicate whether to create the directory or not. Its type signature is as follows:

```
ghci> :t createDirectoryIfMissing
createDirectoryIfMissing :: Bool -> FilePath ->
IO ()
```

If *True* is passed and the directory does not exist, the function will create parent directories as well. If the option is *False*, it will throw up an error:

```
ghci> createDirectoryIfMissing False "/tmp/a/b/c"
*** Exception : /tmp/a/b/c : createDirectory : does
not exist (No such file or directory)
```

```
ghci> createDirectoryIfMissing True "/tmp/a/b/c"
```

```
ghci>
```

You can remove directories using the *removeDirectory* or *removeDirectoryRecursive* functions. Their type signatures are as follows:

```
ghci> :t removeDirectory
removeDirectory :: FilePath -> IO ()
```

```
ghci> :t removeDirectoryRecursive
removeDirectoryRecursive :: FilePath -> IO ()
```

A few examples are shown below:

```
ghci> createDirectoryIfMissing True "/tmp/a/b/c"
ghci>
```

```
ghci> removeDirectory "/tmp/a"
*** Exception : /tmp/a : removeDirectory : unsatisfied
constraints (Directory not empty)
ghci> removeDirectoryRecursive "/tmp/a"
```

```
ghci>
```



The existence of a file can be tested with the *doesFileExist* function. You can check if a directory is present using the *doesDirectoryExist* function. Their type signatures are:

```
ghci> :t doesFileExist
doesFileExist :: FilePath -> IO Bool

ghci> :t doesDirectoryExist
doesDirectoryExist :: FilePath -> IO Bool
```

Some examples that use these functions are shown below:

```
ghci> doesDirectoryExist "/abcd"
False

ghci> doesDirectoryExist "/tmp"
True

ghci> doesFileExist "/etc/resolv.conf"
True

ghci> doesFileExist "/etc/unresolv.conf"
False
```

To know the current directory from where you are running the command, you can use the *getCurrentDirectory* function, and to know the contents in a directory you can use the *getDirectoryContents* function. Their type signatures are:

```
ghci> :t getCurrentDirectory
getCurrentDirectory :: IO FilePath

ghci> :t getDirectoryContents
getDirectoryContents :: FilePath -> IO [FilePath]
```

For example:

```
ghci> getCurrentDirectory
"/tmp"

ghci> getDirectoryContents "/etc/init.d"
["livesys","netconsole",".","..","network","README",
"functions","livesys-late","influxdb"]
```

The *copyFile*, *renameFile* and *removeFile* functions are used to copy, rename and delete files. Their type signatures are shown below:

```
ghci> :t copyFile
copyFile :: FilePath -> FilePath -> IO ()

ghci> :t renameFile
```

```
renameFile :: FilePath -> FilePath -> IO ()
```

```
ghci> :t removeFile
removeFile :: FilePath -> IO ()
```

Here is a very contrived example:

```
import System.Directory
main = do
  copyFile "/etc/resolv.conf" "/tmp/resolv.conf"
  renameFile "/tmp/resolv.conf" "/tmp/resolv.conf.orig"
  removeFile "/tmp/resolv.conf.orig"
```

To obtain the file permissions, use the *getPermissions* function:

```
ghci> :t getPermissions
getPermissions :: FilePath -> IO Permissions

ghci> getPermissions "/etc/resolv.conf"
Permissions {readable = True, writable = False,
executable = False, searchable = False}
```

It is important to separate pure and impure functions in your code and to include the type signatures for readability. An example is shown below:

```
-- Pure
square :: Int-> Int
square x = x * x

-- Impure
main = do
  putStrLn "Enter number to be squared:"
  number <- readLn
  print (square number)
```

The *readLn* function is a parameterised IO action whose type signature is:

```
:t readLn
readLn :: Read a => IO a
```

Executing the code produces the following output:

```
ghci> main
Enter number to be squared :
5
25
```



**By: Shakthi Kannan**

The author is a free software enthusiast and blogs at [shakthimaan.com](http://shakthimaan.com).

# Java Virtual Machine (JVM)

## Delving Deep into its Architecture

A Java Virtual Machine can be thought of as an abstract computer that is defined by certain specifications. The author leads readers deep into the architectural details of JVM to give them a better grasp of its concepts.



A virtual machine, or virtualisation, has emerged as a key concept in operating systems. When it comes to application programming using Java, platform-independent value addition is possible because of its ability to work across different operating systems. The Java Virtual Machine (JVM) plays a central role in making this happen. In this article, let us delve deep into the architectural details of JVM to understand it better.

### Building basics

Let us build our basics by comparing C++ and a Java program with a simple diagram (Figure 1). The C++ compiled object code is OS-specific, say an x86-based Windows machine. During execution, it will require a similar OS, failing which the program will not run as expected. This makes languages like C++ platform- (or OS) dependent. In contrast, Java compilation produces platform-independent byte code, which will get executed using the native JVM. Because of this fundamental difference, Java becomes platform-independent, powered by JVM.

### Exploring JVM architecture

Fundamentally, the JVM is placed above the platform and below the Java application (Figure 2).

Going further down, the JVM architecture pans out as shown in Figure 3. Now let us look into each of the blocks in detail.

In a nutshell, JVM architecture can be divided into two different categories, the details of which are provided below.

**Class loader subsystem:** When the JVM is started, three class loaders are used.

- System class loader:* System class loader maps the class-path environment variables to load the byte code.
- Extension class loader:* Extension class loader loads the byte code from `jre/lib/ext`.
- Bootstrap class loader:* The bootstrap class loader loads the byte code from `jre/lib`.

**Method area:** The method area (or class area) stores the structure of the class once it is loaded by the class loader. The method area is very important; it does two things once the class is stored in this area:

- Identification:* All static members (variable, block, method, etc) are identified from top to bottom.
- Execution:* Static variables and static blocks are executed after the identification phase, and static methods are executed when they are called out. Once all static variables and blocks are executed, only then will the static method be executed.

**Heap area:** The heap area stores the object. Object instances are created in this area. When a class has instance members (instance variable, instance method and instance block), these members are identified and executed only when the instance is created at the heap area.

### Java stacks

In the Java stack area, two threads (main thread and garbage collector) are always running. When the user creates any new thread, it becomes the third thread (Thread-0). When the user creates any method, it is executed by the main thread, inside a stack frame (Figure 4). Each method gets its own stack frame to execute. The stack frame has three sections – the local variable




Figure 1: Differences in C++ and Java compilation

storage section, the instruction storage section and memory slots to perform operations. Each memory slot inside the stack frame is, by default, of 4 bytes, but according to the size of the variable, the size of the slot also shrinks or expands.

According to Figure 4, the addition of a 2-byte variable (a and b) will not result in a byte because the default size of the stack frame memory slot is 4 bytes, which can't be inserted into the byte (r) variable; so we need to typecast it as  $(r = (\text{byte})a+b)$ .

**The PC register:** The program counter (PC) register contains the address of the Java virtual machine instruction currently being executed.

**Native method stacks:** All the native methods are executed in this area.

**Execution engine:** All executions happening in JVM are controlled by the execution engine.

**Native method interface:** Java Native Interface (JNI) enables the Java code running in JVM to call and be called by the native application and libraries (Native Method Libraries) written in other languages such as C and C++.

## JVM in action

Now let us take a look at a few Java code snippets and understand the role of various JVM components, during execution.

**Example 1 (when all members are static):** To understand this example, the method area is explained earlier in this article. According to the method area, all static members of a class are identified and executed in the same order in which they appear. When all static members are executed, only then is the main function executed.

```
class Test{
    static int a =m1();
    static{
        System.out.println ("in static block");
    }
    public static int m1(){
        System.out.println("in m1");
        return 10;
    }
    public static void main(String[] args){
        System.out.println("in main");
    }
}
```

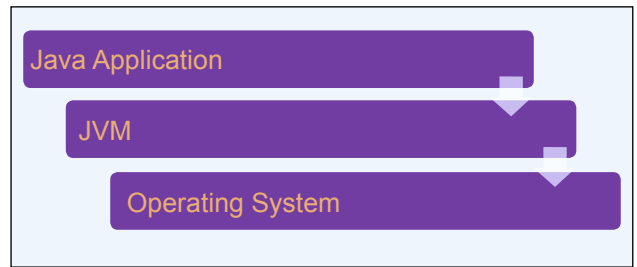


Figure 2: How JVM fits between the OS and Java application

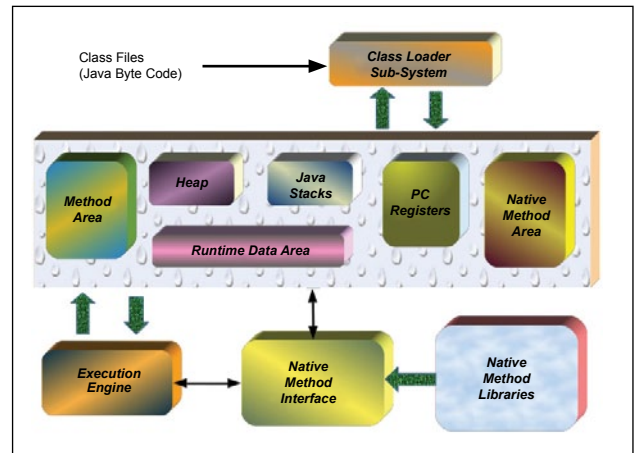


Figure 3: Architecture of JVM

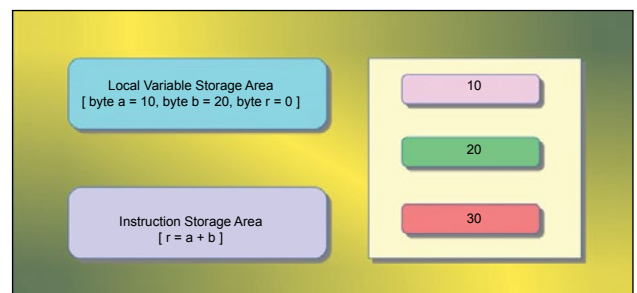


Figure 4: Architecture of the stack frame

```
}
The output of the above program is:
in m1
in static block
in main
```

**Example 2 (in case of inheritance):** To understand this example, you need to understand the method area as well as Example 1. In this example, we are trying to explain that, when a class is inherited from any other class, the static members are identified and executed from the top to the bottom or from parent class to child class, and the main function will be executed at the end.

```
class Vehicle{
    static int a = m1();
    public static int m1(){
```



```

        System.out.println("in m1");
        return 10;
    }
    static{
        System.out.println("Vehicle static block");
    }
}
class Car extends Vehicle{

    static int b = m2();
    public static int m2(){
        System.out.println("in m2");
        return 10;
    }
    static{
        System.out.println("Car static block");
    }
    public static void main(String[] args){
        System.out.println("in main");
    }
}

```

The output of the above code is:

```

in m1
Vehicle static block
In m2
Car static block
in main

```

The above output clarifies one thing—that the order of identification and the execution of static members in case of inheritance will occur from top to bottom or from parent class to child class.

**Example 3 (when members are static, non-static and constructor):** In this example, the effect of the heap area on the Java program is explained (the heap area itself has been covered earlier in the article). All the Java objects and instance members are initialised here. The method area is about static members and the heap area is about object and non-static members. In the heap area, when the Java objects are created, only then is the instance variable identified and executed. The constructors are executed in the end.

```

class Test{
    static int a=m1();
    int b=m2();
    {
        System.out.println("instance block");
    }
    public int m2(){
        System.out.println("in m2");
        return 10;
    }
}

```

```

static {
    System.out.println("in static block");
}
public static int m1(){
    System.out.println("in m1");
    return 15;
}
public Test(){
    System.out.println("in constructor");
}
public static void main(String[]args){
    System.out.println("in main");
    Test t = new Test();
}
}

```

The output of the above code is:


```

in m1
in static block
in main
in m2
instance block
in constructor

```

The above output clarifies the following three points:

1. The static member is both identified as well as executed first.
2. Instance members (non-static variables and non-static methods) are identified and executed only if the instance is created, and the order of identification and execution will be from top to bottom.
3. Constructors are always executed last.

The Java programming language becomes platformindependent because of JVM and the nature of its byte code. Because of this, Java has scaled across multiple machines, platforms and devices, powering enterprise Web applications as well as consumer mobile applications. The architecture of JVM has a significant effect on a Java program. Both static and non-static members of a Java class are treated differently by JVM (examples 1 and 3). Static members of a class are identified and executed as soon as the class is loaded. Non-static members are identified and executed only if the instance of the class is created. 

### By: Vikas Kumar Gautam

The author is a mentor at Emertxe Information Technology (P) Ltd. His main areas of expertise include application development using Java/J2EE and Android for both Web and mobile devices. A Sun Certified Java Professional (SCJP), his interests include acquiring greater expertise in the application space by learning from the latest happenings in the industry. He can be reached at [vikash\\_kumar@emertxe.com](mailto:vikash_kumar@emertxe.com)

# Developing Applications Using NoSQL Databases



NoSQL (or 'Not only SQL') databases offer the means to store and retrieve data that is not stored in the traditional RDBMS style of tabular columns. NoSQL is rapidly finding a place in Big Data and realtime Web applications.

For the past few years, the world of Web technologies has been associated with assorted programming languages and scripts. The domain of Web development is not confined to a specific programming language or library, using which Web applications and portals are developed. Thousands of toolkits, programming paradigms, scripts, databases and application programming interfaces (APIs) are in use for multiple services and applications. The days when only a static website was needed using classical hypertext markup language (HTML) are long gone. The corporate world is using different tools to design, develop and launch applications with maximum user interactions as well as effective graphical user interfaces (GUI).

## SQL, NewSQL and NoSQL databases

Whenever there is need of a Web 2.0 portal (an interactive website), database-oriented applications are required, as the back-end database keeps and maintains the records required for the appropriate functioning of the modules. For example, guest book messages, posts, blogs, e-mail messages, chat messages and comments are stored in the back-end databases so that these can be retrieved, processed or edited at any instant.

Classically, the RDBMS packages used for database applications include MySQL, Oracle, Apache Derby, IBM DB2, IBM Notes, Microsoft SQL Server, PostgreSQL, SQLite, Sybase and many others. These are known as the traditional

SQL databases, which are ACID properties compliant.

NewSQL is a contemporary relational database management system, which provides the same scalable performance as NoSQL systems for online transaction processing (OLTP) read-write workloads, and maintains the ACID guarantees of a classical database system.

Nowadays, Web applications use data in heterogeneous formats, which includes audio, video, text, streaming content, signals, images, pixels and many others. In each file, there are a number of file formats. For example, in video, there are a number of file formats including MPEG, MP4, AVI, 3GP, WMV, OGG, FLV and others. In the same manner, image or graphics file formats include GIF, PNG, JPEG, PCX, BMP, TIFF and many others.

Now the major issue is the compatibility of the Web application with all these file formats in different domains. It is at this point that implementing NoSQL databases makes the most sense. In NoSQL databases, any type of file format can be processed and integrated in the Web application.

NoSQL databases provide a storage and retrieval system that is different from the tabular relations used in relational databases. The data structure in NoSQL databases is entirely different from the classical RDBMS. NoSQL databases are rapidly being used in Big Data and realtime Web applications.

There have been various approaches to classifying NoSQL databases, each with different categories and sub-categories. Because of the variety of approaches and the way they overlap, it is difficult to get a clear overview of non-relational databases.

#### NoSQL Databases for Multiple Domains

**Column:** Accumulo, Cassandra, Druid, HBase

**Document:** Clusterpoint, Apache CouchDB, Couchbase, MarkLogic, MongoDB

**Key-Value:** Dynamo, FoundationDB, MemcacheDB, Redis, Riak, FairComc-treeACE, Aerospike

**Graph:** Allegro, Neo4J, InfiniteGraph, OrientDB, Virtuoso, Stardog

MongoDB is one of the prominent cross-platform document-oriented NoSQL databases released under a combination of the GNU Affero GPL and the Apache Licence. It is free and open source software.

According to media reports, the database technology used for the back-end of the Aadhaar cards in India is MongoDB. Aadhaar implementation involves enrolling thousands of Indians each day and adds terabytes of data to the data repository. MongoDB as well as other data management and analytics software providers continue to produce insights that aid Aadhaar and the lives of millions of India's citizens.

MongoDB's customers include companies such as Disney, the New York Times, Cisco, MTV, Forbes, Craigslist, ADP, AstraZeneca, FourSquare, IBM, Intuit, Microsoft, McAfee, UnderArmour, and eBay.

Initially developed by 10gen (now MongoDB Inc) in

October 2007 as a component of a planned platform (as a service product), the organisation shifted to an open source development model in 2009, with 10gen offering commercial support and other services. Since then, MongoDB has been adopted by a number of major websites and services.

## Web application development using PHP and NoSQL databases

In order to use MongoDB with PHP, we need the MongoDB PHP driver.

Download the driver from the URL *Download PHP Driver*.

Now unzip the archive and use *php\_mongo.dll* in your PHP extension directory ("ext" by default) and add the following line to your *php.ini* file:

```
extension=php_mongo.dll
```

## Database connection and selecting the database

To make a connection, we need to specify the database's name. If it doesn't exist, then MongoDB creates it automatically.

```
<?php
// connect to mongodb
$m = new MongoClient();
echo "Connection to database successfully";
// select a database
$db = $m->mydb;
echo "Database mydb selected";
?>
```

When the program is executed, it will give the following result:

```
Connection to database successfully
Database mydb selected
```

## Creating a collection

To create a collection, type:

```
<?php
// connect to mongodb
$m = new MongoClient();
echo "Connection to database successfully";
// select a database
$db = $m->mydb;
echo "Database mydb selected";
$collection = $db->createCollection("mycol");
echo "Collection created successfully";
?>
```



When the above program is run, it will give the following result:

```
Connection to database successfully
Database mydb selected
Collection created successfully
```

## Inserting a document

To insert a document into MongoDB, the *insert()* method is used:

```
<?php
// connect to mongodb
$m = new MongoClient();
echo "Connection to database successfully";
// select a database
$db = $m->mydb;
echo "Database mydb selected";
$collection = $db->mycol;
echo "Collection selected successfully";
$document = array(
    "title" => "MongoDB",
    "description" => "database",
    "likes" => 100,
    "url" => "http://www.mynosqlldb.com/mongodb/",
    "by", "My NoSQL Implementation"
);
$collection->insert($document);
echo "Document inserted successfully";
?>
```

The above program will give the following result:

```
Connection to database successfully
Database mydb selected
Collection selected successfully
Document inserted successfully
```

## Finding all documents

To select all documents from the collection, the *find()* method is used:

```
<?php
// connect to mongodb
$m = new MongoClient();
echo "Connection to database successfully";
// select a database
$db = $m->mydb;
echo "Database mydb selected";
$collection = $db->mycol;
echo "Collection selected successfully";

$cursor = $collection->find();
```

```
// iterate cursor to display title of documents
foreach ($cursor as $document) {
    echo $document["title"] . "\n";
}
?>
```

When the above program is executed, it will display the following result:

```
Connection to database successfully
Database mydb selected
Collection selected successfully
{
    "title": "MongoDB"
}
```

## Updating a document

To update a document, you need to use the *update()* method:

```
<?php
// connect to mongodb
$m = new MongoClient();
echo "Connection to database successfully";
// select a database
$db = $m->mydb;
echo "Database mydb selected";
$collection = $db->mycol;
echo "Collection selected successfully";
// now update the document
$collection->update(array("title"=>"MongoDB"), array('$set'=>array("title"=>"MongoDB Implementation")));
echo "Document updated successfully";
// now display the updated document
$cursor = $collection->find();
// iterate cursor to display title of documents
echo "Updated document";
foreach ($cursor as $document) {
    echo $document["title"] . "\n";
}
?>
```

When the above program is executed, it will give the following result:

```
Connection to database successfully
Database mydb selected
Collection selected successfully
Document updated successfully
Updated document
{
    "title": "MongoDB Implementation"
}
```

...Continued on page 64



# The Need to Address Security Concerns in the Cloud Environment

Security of data is of paramount concern to those in government and industry. This article takes readers through the typical security concerns in the traditional environment, before going on to the cloud environment.

**B**efore we discuss security concerns in the cloud, let us take a small detour to look at the security situation in traditional environments. In today's competitive market, a typical IT environment has fewer capabilities to handle dynamic requirements due to budget constraints and a large user base. To survive in the current state and to plan for long-term benefits, organisations need to reduce costs, increase storage capacity, as well as automate the monitoring of IT environments and the management of resources for disaster recovery and availability, so as to achieve a flexible IT environment.

The following are the challenges in the traditional environment:

- Security systems utilise excessive amounts of resources
- Vulnerability scanning or monitoring tools may not detect vulnerability threats due to the fast growing volume of data
- Accurate and timely information about threats is not widely shared; information about the vulnerabilities remains with a smaller group and not all the stakeholders
- Audit systems are not built to operate on the dynamic nature of data centres
- Frequent failovers occur due to limitations on scalability
- Aggressive cost-cutting plans are degrading security programs at a time when threats are escalating
- Security access privileges and roles are not as per standards

Complexities of the underlying infrastructure in on-premise environments have attracted organisations to the

cloud environment to achieve agility, high availability and reliability. Thus, an organisation's security and compliance requirements also need to be aligned and automated in an agile cloud environment to realise the desired benefits of moving to the cloud.

## Security in the cloud

Cloud computing provides the next generation of Internet-based, scalable and distributed computing systems in which resources are offered 'as a service'. Today's IT organisations are under increasing pressure to securely deploy applications in public, private or hybrid clouds. Security in the cloud environment is a major concern. Despite the likely gains achieved from cloud computing, the security model is still not entirely fool-proof and hence it impacts cloud adoption. Traditional security controls and policies are not the best fit for the virtual world. Multi-tenancy, virtualisation and geographical differences have made the boundaries of the security model more dynamic than before. Cloud security is no longer the 'elephant in the room' since various organisations and standards bodies have raised their concerns about it. It is the general public as well as the providers and consumers of cloud services who seem to continue in a state of ignorant bliss. Cloud consumers and service providers must realise that a stitch in time will save nine.

Infrastructure abstraction and lack of visibility in cloud resources can create a number of security and compliance issues. In addition to the usual challenges of developing

secure IT systems, cloud computing presents an added level of risk because:

- Services are outsourced to a third party
- It requires moving an application into hostile territory
- Multi-tenancy
- Loss of governance
- Legal and contractual risks

Cloud security is basically about three goals or objectives:

- Confidentiality (C)
- Integrity (I)
- Availability (A)

Cloud security and its compliance are the key components that are needed to protect cloud infrastructure against ever evolving new threats. It helps if organisations rationalise compliance requirements, controls, standards and best practices into centralised security policies administered consistently across virtual and physical infrastructures.

There are certain key areas where security levels are likely to be questionable and addressing them in an efficient manner is a critical task for easy and swift adoption of the cloud environment. Let's look at these areas more closely.

### Identity and access management (IAM)

Unauthorised access to information resources in the cloud is a primary concern. In the current state of IAM technology, standards supported by cloud services providers offering various cloud service models (such as SaaS, PaaS and IaaS) are not consistent across providers. Google, Microsoft and Salesforce.com seem to demonstrate basic IAM capabilities. IAM should include the following: identity provisioning and de-provisioning, information privacy, linking, mapping, federation, attributes federation, authentication, SSO and authorisation.

### Data security

Confidentiality, integrity and availability (CIA) form the golden trio of data security. Data security becomes more important while using cloud computing for all service models: IaaS, PaaS, and SaaS. Cloud environments are shared with multiple tenants and cloud service providers have privileged access to the data in those environments. Sensitive data stored in a cloud environment must be secured using a combination of strong encryption, access control, contractual liability, etc.

### Network security

At the network level of infrastructure security, it is critical to distinguish between public clouds and private clouds. With private clouds, there may not be new attacks, vulnerabilities or changes in risk that information security personnel need to consider. In public clouds, changing security requirements will warrant more attention considering that it is the cloud service provider that is in control of resources.




Figure 1: Cloud security concerns



Figure 2: C-I-A

### Virtualisation security

Virtualisation technologies enable multi-tenancy in cloud business models by providing a scalable and shared resource platform for all tenants. The threat of an unauthorised virtual machine (VM) is far higher because it is pretty easy to create and deploy virtual machines. Organisations need to enforce strict policies on the use of such virtual machine environments. VM images can be copied, along with the data and applications that they hold. These images can be brought back online to an unsecured network. It becomes easier for an attacker to access the contents managed within the copied image. Deployment of virtualisation in any organisation should be a controlled and policy-driven roll-out like any other computer platform, software or application.

### Physical security

Cloud resources such as servers, routers, storage devices and power supplies that support cloud operations should be physically secure. Safeguards include the adequate control and monitoring of physical access using biometric access control measures and closed circuit television (CCTV) monitoring. Cloud service providers need to clearly explain how they manage physical access to the servers that host client workloads and support client data.

### Use cases

#### Application firewall to protect apps in the private cloud

Let's look at a hypothetical case. The XYZ Market is a new shopping channel led by two seasoned industrialists. They have a private cloud for their organisation. XYZ Market's HR and business departments have applications named 'Time and Expense management' and 'E-commerce' respectively, which they want to deploy on a private cloud environment to gain agility, high availability and fault tolerance. Each application has two tiers – Web and the data base. The main concern of both departments is how to keep both applications secure in a private cloud environment.

The departments need to enforce the following separation between the applications and tiers of each application using the networking and security firewall



| Cloud security concerns        |  |
|--------------------------------|--|
| Category                       | Concern  |
| Identity and access management | Authentication   |
|                                | SSO  |
|                                | Authorisation  |
|                                | User directory and federation services                             |
| Data security                  | Data confidentiality   |
|                                | Data integrity   |
|                                | Availability   |
|                                | Back-ups and archives  |
|                                | Key management   |
| Network security               | Security for data in transit                                       |
|                                | Perimeter security   |
|                                | Network security threats (DoS, man-in-the-middle, packet sniffing) |
| OS and virtualisation security | Virtualisation/Hypervisor threats                                  |
|                                | Host OS  |
|                                | OS hardening   |
|                                | Anti-virus   |

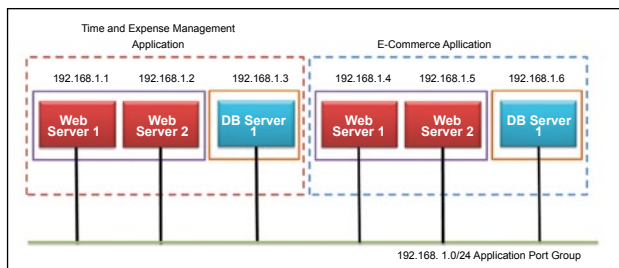


Figure 3: Application deployment on a private cloud

provided by specific private cloud products.

- Ensure complete isolation between Application 1 and Application 2
- Isolate one Web server from another in each application
- Allow HTTP (80)/HTTPS (443) traffic to Web servers
- Allow Web server to database server communication on a specific DB port
- Block all other traffic
- Security policies to be completely transparent to IP address changes and network renumbering
- Dynamic rules to be applied to new virtual machines automatically
- Default 'Deny All' policy in case of firewall rules

## The solution

This solution gives organisations deep visibility into network communications between virtual machines, and eliminates the hardware and policy sprawl associated with using traditional appliances to secure virtualised

## Best practices to ensure security in a cloud environment

- Design for failure
- Geo-distributed data centre
- Back-up of data on the cloud
  - Full back-up
  - Incremental back-up
- Layer-wise security
  - Web layer
  - App layer
  - Data layer
- Default 'Deny all' policy
- Logical data centre
- Geo-distributed app architecture
- Software-based data replication
- Data durability
- Use of Memcache
- Monitoring services
- Distribute static and dynamic content

applications. In addition to that, it optimises hardware resource utilisation while maintaining strong levels of security.

Security offerings from proprietary and open source initiatives enable a hypervisor-based application-aware firewall to dynamically define and enforce policies for virtual machines that are supporting business critical applications. Customers can define a security group, for instance, to protect a group of virtual machines that might hold credit card data that needs to be PCI compliant, or another group of virtual machines that hold patient health records and needs to be HIPAA compliant.

Firewall rules to accomplish the requirements of this use case must be defined. When a new virtual machine joins the container or security group, the rules set-up will be applied automatically and there will be no need to define new rules.

## Advantages

- Eliminates the complexity and cost of using security groups to segregate traffic. An application firewall with deep visibility helps to filter traffic between virtual machines.
- Monitors 'flows' of traffic, independent from firewall rules.
- Reports are based on applications, sessions, bytes and other details, allowing administrators to define and refine their policies based on actual traffic patterns.

## Identity management

In an organisation, multiple business units deploy their applications in the private cloud. In addition to that, customers also use private cloud resources. The challenge is to provide secure access to all consumers of a private cloud. The security and identity infrastructure in the private cloud is an important management platform component. Private clouds can be based on VMware or OpenStack products.

## The solution

Multiple directory services are used in organisations such as OpenLDAP, which is a free, open source

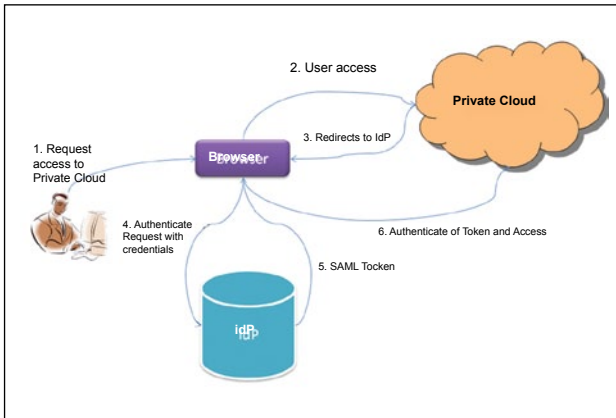



Figure 4: Identity management

implementation of the Lightweight Directory Access Protocol (LDAP), developed by the OpenLDAP Project. Active Directory (or AD) is a directory service created by Microsoft for Windows domain networks and is included in most Windows Server operating systems.

Administrators can set up a separate identity provider (IdP) using either OpenAM, ADFS or any other open source IAM product. The use of the cloud manager component provided by private cloud products helps XYZ Market set the OpenAM IdP as the identity provider for it. Thereafter, when users from XYZ log in to their organisation's private cloud, they will be redirected

to OpenAM, where they can be authenticated and be directed back to the cloud portal.

## Advantages

- Secure access to private cloud infrastructure
- No need for XYZ to create its own user identities in order to participate in the management process, which would increase the administrative overhead **END** 

## Resources

- [1] <http://www.cloudsecurityalliance.org/cm.html>
- [2] <http://www.computing.co.uk/computing/comment/2270462/securing-public-private-clouds>
- [3] <http://blog.securityarchitecture.com/2010/08/major-cloud-computing-privacy-legal.html>
- [4] <http://www.opensecurityarchitecture.org/cms/library/patternlandscape/251-pattern-cloud-computing>
- [5] <http://csrc.nist.gov/publications/nistpubs/800-27A/SP800-27-RevA.pdf>
- [6] <http://searchcloudsecurity.techtarget.com/tip/Cloud-security-services-WAFs-and-DdoS-attack-prevention>
- [7] <http://basho.com/tag/ecommerce/>
- [8] <http://www.3dcart.com/pci-compliance.html>

## By: Mitesh Soni

The author is a technical lead at IGATE. He is in the cloud practices group and loves to explore new technologies. Blog: <http://clean-clouds.com>

...Continued from page 60

## Deleting a document

To delete a document, you need to use the `remove()` method:


```
<?php
// connect to mongodb
$m = new MongoClient();
echo "Connection to database successfully";
// select a database
$db = $m->mydb;
echo "Database mydb selected";
$collection = $db->mycol;
echo "Collection selected successfully";

// now remove the document
$collection->remove(array("title"=>"MongoDB
Implementation"),false);
echo "Documents deleted successfully";

// now display the available documents
$cursor = $collection->find();
// iterate cursor to display title of documents
echo "Updated document";
```

```
foreach ($cursor as $document) {
    echo $document["title"] . "\n";
}
?>
```

When the above program is executed, it will produce the following result:

```
Connection to database successfully
Database mydb selected
Collection selected successfully
Documents deleted successfully END 
```


## By: Dr Gaurav Kumar and Amit Doegar

Dr Gaurav Kumar is the MD of Magma Research & Consultancy Pvt Ltd, Ambala. He is associated with a number of academic institutes, delivering expert lectures and conducting technical workshops on the latest technologies and tools. E-mail: [kumargaurav.in@gmail.com](mailto:kumargaurav.in@gmail.com)

Amit Doegar is assistant professor in the National Institute of Technical Teachers' Training and Research at Chandigarh. He can be contacted at [amit@nittrchd.ac.in](mailto:amit@nittrchd.ac.in)

# Automate the Bare Metal Provisioning Process through Razor

This article describes how to automate the bare metal provisioning of physical and virtual machines through Razor, an open source tool that works perfectly with Puppet.



**R**azor was created to automatically discover bare metal hardware and dynamically configure operating systems and hypervisors. Razor makes it easy to provision a node with no previously installed operating system and bring it under the management of Puppet. It was originally developed by EMC and is based on Tiny Core Linux. The Razor micro-kernel is 64-bit only. Razor can only provision 64-bit machines. It has the ability to discover hardware via in-memory instances of the Razor micro-kernel (a.k.a. Razor MK). The source code of the micro kernel is available at: <https://github.com/puppetlabs/razor-el-mk> under GPL v2 license.

Razor is completely open source, which means that it offers you the freedom to build your own custom Razor MK images, with the option to specify user accounts, the ability to enable remote SSH access for debugging, and to build and include custom Tiny Core Linux extensions to support unique hardware for your environment. Razor's policy-based bare-metal provisioning lets you make an inventory and manage the lifecycle of your physical machines.

## How does Razor work?

Whenever a new node gets added, Razor discovers its

characteristics by booting it with the Razor micro-kernel and inventorying its facts. The node is tagged based on its characteristics. Tags contain a match condition—a Boolean expression that has access to the node's facts, and determines whether the tag should be applied to the node or not. Node tags are compared to tags in the policy table. The first policy with tags that match the node's tags is applied to the node.

## Provisioning elements of Razor

**Repositories:** These take care of 'What to install?' They basically indicate the contents to be installed on a system. To create a repository, either import or install an ISO or point at an existing package repository.

**Tasks:** A task takes care of 'How to install' using installation scripts such as kickstart files, preseed files and additional shell scripts. Predefined tasks are shipped with Razor, and custom tasks can easily be added without additional coding.

**Broker:** This takes care of 'How to manage' with post-installation scripts that install a configuration management agent on the node and enrol the node with the configuration management system (e.g., Puppet).

**Tag:** This takes care of 'Where to install' with the Boolean



expression that uses node facts and metadata. Tags are used to match nodes and policies.

**Policy:** This takes care of ‘Combining it all’ with the use of ordered tables which combine all the above elements in the form of YAML.

## Setting up the Razor server

It is recommended that the Razor server shouldn’t be installed on the same machine on which the Puppet master is running. The reason is that the default port for Razor is 8080, which conflicts with the default Puppet DB port. To set up a test environment, we will need at least 2 VMs – one for the Puppet master and the other for the Razor server:

- Puppet server (hostname - *puppetmaster*)
- Razor server (hostname - *puppetagent1*)

Razor has been specifically validated on RHEL/CentOS 6.5 but it should work on all 6.x versions. I assume that the Puppet server is installed and configured properly on CentOS 6.5 VM with the hostname *puppetmaster*. If you are new to Puppet, I recommend reading [https://docs.puppetlabs.com/guides/install\\_puppet/install\\_el.html](https://docs.puppetlabs.com/guides/install_puppet/install_el.html)

Here are the steps that need to be followed to set up a Razor server on the *puppetagent1* machine.

**Installing a Razor module:** A Razor module is available under the Puppet Labs GitHub repository. I assume that the Git package is already installed on *puppetagent1*. We will use Rubygems software (rightly called a ‘gem’) which allows you to easily download, install and use Ruby packages on the system.

```
# gem install bundler
# cd /opt; git clone https://github.com/puppetlabs/razor-server.git
# cd razor-server;
# bundle install
# rake db:migrate
# torquebox deploy
#yum install jruby
#curl -L -O http://torquebox.org/release/org/torquebox/torquebox-dist/3.0.1/torquebox-dist-3.0.1-bin.zip
#unzip torquebox-dist-3.0.1-bin.zip -d $HOME
# jruby bin/razor-admin -e production migrate-database
```

**Set the following environmental variable:**

```
#cat /root/.bashrc
export TORQUEBOX_HOME=$HOME/torquebox-3.0.1
export JBOSS_HOME=$TORQUEBOX_HOME/jboss
export JRUBY_HOME=$TORQUEBOX_HOME/jruby
```




Figure 1: Razor discovering capability

```
[root@puppetagent1 ~]# su - postgres
-bash-4.1$ cd
-bash-4.1$ psql -h 127.0.0.1 -l -U razor
Password for user razor:
List of databases
Name | Owner | Encoding | Collate | Ctype | Access privilege
-----+-----+-----+-----+-----+-----
postgres | postgres | UTF8 | en_US.UTF-8 | en_US.UTF-8 |
razor | razor | UTF8 | en_US.UTF-8 | en_US.UTF-8 |
razor_dev | razor | UTF8 | en_US.UTF-8 | en_US.UTF-8 |
razor_prd | razor | UTF8 | en_US.UTF-8 | en_US.UTF-8 |
razor_test | razor | UTF8 | en_US.UTF-8 | en_US.UTF-8 |
template0 | postgres | UTF8 | en_US.UTF-8 | en_US.UTF-8 | =c/postgres
+ | | | | | postgres=C/r/pos
template1 | postgres | UTF8 | en_US.UTF-8 | en_US.UTF-8 | =c/postgres
+ | | | | | postgres=C/r/pos
postgres (7 rows)
```

Figure 2: Verifying the Razor database connectivity

```
export PATH=$JRUBY_HOME/bin:$PATH
```

**Installing the database:** Razor uses PostgreSQL as its database server. To configure the database, follow the steps:

```
# yum remove postgresql postgresql-server
# curl -O http://yum.postgresql.org/9.4/redhat/rhel-6-x86_64/pgdg-centos94-9.4-1.noarch.rpm
# rpm -ivh pgdg-centos94-9.4-1.noarch.rpm
# service postgresql-9.4 initdb
# chkconfig postgresql-9.4 on
Log in as psql user and verify the table entry (see Figure 2).
```

**Installing the micro-kernel:** Download a pre-built micro-kernel from <http://links.puppetlabs.com/razor-microkernel-latest.tar>. The micro-kernel is based on Fedora 19 and needs to be manually put into your *repo\_store\_root* directory; it cannot be added using the API. If you downloaded the prebuilt micro-kernel above, simply extract it into your *repo\_store\_root* directory. Doing so will create a sub-directory called *microkernel* with its contents.

```
#cd /var/lib/razor/repo-store
# wget http://links.puppetlabs.com/razor-microkernel-latest.tar
# tar xvf razor-microkernel-latest.tar
#cd microkernel
# ls
initrd0.img README SHA256SUM SHA256SUM.sig vmlinuz0
```

**Configuring the database:** Edit */opt/razor/config.yaml* and change the database URL setting. Once that is done, you can load the Razor database schema into your PostgreSQL database, and finally start the service (see Figure 3).

Ensure that you have the following line: *repo\_store\_root: /var/lib/razor/repo-store* placed under */opt/razor/config.yaml*. Verify that razor-server service is in a running state:

```
#service razor-server status
razor-server is running (pid 1380)
```

```
production:
  database_url: 'jdbc:postgresql:razor_prd?user=razor&password=razor'
development:
  database_url: 'jdbc:postgresql:razor_dev'
test:
  database_url: 'jdbc:postgresql:razor?user=razor&password=razor'
all:
```

Figure 3: Razor configuration file

```
[root@puppetagent1 tftpbboot]# pwd
/tftpbboot
[root@puppetagent1 tftpbboot]# ls -la
total 80
drwxr-xr-x. 2 root root 4096 Jun 19 07:21 .
dr-xr-xr-x. 24 root root 4096 Jun 19 07:50 ..
-rw-r-xr-x. 1 root root 1072 Jun 19 07:21 bootstrap.ipxe
-rw-r-xr-x. 1 root root 66979 Jun 19 07:21 undionly.kpxe
```

Figure 4: Razor iPXE bootstrap script

```
IPXE 1.0.0* (53653) -- Open Source Network Boot Firmware -- http://ipxe.org
Features: HTTP ISCSI DNS TFTP AoE SRP bzImage ELF MBOOT PXE PXEXT Menu

net0: 00:0c:29:04:a1:a3 using 02545em on PCI02:00.0 (open)
(LinkUp: TX:0 TXE:0 RX:0 RXE:0)
Configuring (net0 00:0c:29:04:a1:a3)..... Error 0x040ee119 (http://
ipxe.org/040ee119)
net1: 00:0c:29:04:a1:a3 using 02545em on PCI02:01.0 (open)
(LinkUp: TX:0 TXE:0 RX:0 RXE:0)
Configuring (net1 00:0c:29:04:a1:a3).... ok
net1: 192.168.1.110/255.255.255.0 gw 192.168.1.55
Next server: 192.168.1.55
Filename: bootstrap.ipxe
tftp://192.168.1.55/bootstrap.ipxe... ok
chain booting the razor script from puppetagent1.cse.com
http://puppetagent1.cse.com:8080/svc/boot?net0=00-0c-29-04-a1-a3&net1=00-0c-29-0
4-a1-a3&net2=00-0c-29-04-a1-a3&serial=0&mac=563204d220dd206f2033320e4220a1k200
c-1420442201d2208c207032084220a1k20a3&asset=No208asset20Tag&uid=564dd6f-33e
4-a10c-144d-1d8c7084a1a3..... ok
http://puppetagent1.cse.com:8080/svc/repo/microkernel/initrd0.img... 0%_
```

Figure 5: Razor micro-kernel booting and inventorying its facts

```
IPXE 1.0.0* (7485) -- Open Source Network Boot Firmware -- http://ipxe.org
Features: HTTP ISCSI DNS TFTP AoE bzImage ELF MBOOT PXE PXEXT Menu

net0: 00:0c:29:04:a1:a3 using undionly on UNDI-PCI02:01.0 (open)
(LinkUp: TX:0 TXE:0 RX:0 RXE:0)
DHCP (net0 00:0c:29:04:a1:a3).... ok
net0: 192.168.1.110/255.255.255.0 gw 192.168.1.55
Next server: 192.168.1.55
Filename: bootstrap.ipxe
tftp://192.168.1.55/bootstrap.ipxe... ok
chain booting the razor script from puppetagent1.cse.com
http://puppetagent1.cse.com:8080/svc/boot?net0=00-0c-29-04-a1-a3&net1=00-0c-29-0
4-a1-a3&net2=00-0c-29-04-a1-a3&serial=0&mac=563204d220dd206f2033320e4220a1k200
c-1420442201d2208c207032084220a1k20a3&asset=No208asset20Tag&uid=564dd6f-33e
4-a10c-144d-1d8c7084a1a3... ok
http://puppetagent1.cse.com:8080/svc/repo/microkernel/initrd0.img... ok
http://puppetagent1.cse.com:8080/svc/repo/microkernel/mlinuz0... ok
```

Figure 6: Bootloader loading the kernel image

```
DHCP request via eth0
^C
[root@puppetagent1 tftpbboot]# ^C
[root@puppetagent1 tftpbboot]# cd
[root@puppetagent1 ~]# cd /etc/puppetlabs/
[root@puppetagent1 puppetlabs]# ls
installers mcollective puppet
[root@puppetagent1 puppetlabs]# cd puppet/
[root@puppetagent1 puppet]# ls
auth.conf hiera.yaml puppet.conf ssl
[root@puppetagent1 puppet]# cd
[root@puppetagent1 ~]# find / -name microkern
[root@puppetagent1 ~]# razor nodes
From http://localhost:8080/api/collections/nodes:

+-----+-----+-----+-----+-----+
| name | dhcp_mac | tags | policy | metadata count |
+-----+-----+-----+-----+-----+

```

Figure 7: Status of Razor nodes before provisioning

After you've followed one of the above installation methods, you should be able to go to <http://localhost:8080/api> and get the API entry point that will give you a JSON document that talks about *collections* and *commands*.

**Setting up PXE:** Type the following commands (see Figure 4).

```
# wget http://links.puppetlabs.com/pe-razor-ipxe-firmware-3.3.
# cp undionly-20140116.kpxe /var/lib/tftpbboot
# cp bootstrap.ipxe /var/lib/tftpbboot
```

**Configuring DNSMASQ:** Set the following configuration under `/etc/dnsmasq.conf`.

```
# This works for dnsmasq 2.45
```

```
# ipxe sets option 175, mark it for network IPXEBOOT
dhcp-match=IPXEBOOT,175
dhcp-boot=net:IPXEBOOT,bootstrap.ipxe
dhcp-boot=undionly.kpxe
# TFTP setup
enable-tftp
tftp-root=/var/lib/tftpbboot
dhcp-range=192.168.1.50,192.168.1.150,12h
```

This completes the Razor server configuration. Now let's create a new VM and try to PXE boot.

As you've seen above, the new VM listened to the net1 interface and acquired the IP address from the DHCP server. Next, Razor discovers its characteristics by booting it with the Razor micro-kernel and inventorying its facts. Meanwhile, you can check the status of nodes as shown in Figure 7.

As seen in this figure, as of now, there are no provisioning elements created for the new PXE booted VM. It's time to create the provisioning elements of Razor. I have created a provisioning element for CentOS 6.5 x64.

## Creating the repository

To create a repository, type:

```
[root@puppetagent1 ~]# razor create-
repo --name=CentOS6.5-Repo --iso-url http://
192.168.1.100/OS/Linux/CentOS/CentOS-6.5-x86_64-bin-DVD1.iso
--task centos6
From http://localhost:8080/api:
name: CentOS6.5-Repo
iso_url: http://192.168.1.100/OS/Linux/CentOS/CentOS-
6.5-x86_64-bin-DVD1.iso
url: ---
task: ---
command: http://localhost:8080/api/collections/commands/1
```

## Creating a broker

To create a broker, type:

```
[root@puppetagent1 ~]# razor create-
broker --name foo --broker-type puppet-pe --
configuration '{ "server": "puppetmaster.cse.com" }'
From http://localhost:8080/api:
name: foo
broker-type: puppet-pe
configuration:
server: puppetmaster.cse.com
policies: 0
command: http://localhost:8080/api/collections/
commands/2
```

## Creating policy

To create a policy, type:

```
[root@puppetagent1 ~]# cat policy.json
{
    "name": "test_node1",
    "repo": "Centos-6.4",
    "task": "centos",
    "broker": "pe",
    "enabled": true,
    "hostname": "node${id}.cse.com",
    "root_password": "razor123",
    "max_count": 20,
    "tags": ["small"]
}
```

We are good to create a policy through the following command:

```
[root@puppetagent1 centos.task]# razor create-policy --name
demo --repo CentOS6.5-Repo --hostname 'host${id}.cse.com'
--root-password 'dell01' --broker foo --tag test --task
centos
From http://localhost:8080/api:
    name: demo
    repo: CentOS6.5-Repo
    task: centos
    broker: foo
    enabled: true
    max_count: nil
    tags: test
    nodes: 0
    command: http://localhost:8080/api/collections/commands/6
```

## Creating tasks

By default, Razor is shipped with CentOS tasks under `/opt/razor/tasks/centos.task`.

Once you successfully create the policy, you can see the following output:

```
# razor nodes
From http://localhost:8080/api/collections/nodes:

+-----+-----+-----+-----+-----+
| name | dhcp_mac | tags | policy | metadata count |
+-----+-----+-----+-----+-----+
| node1 | 00:0c:29:04:a1:ad | test | --- | 0 |
+-----+-----+-----+-----+-----+
```

Query an entry by including its name, e.g., ``razor nodes node1``

To get the detailed information about the node1, run the following command:

```
#razor nodes <node name> facts
```

Finally, after a few minutes, a new CentOS 6.5 VM comes up.

```
[root@puppetagent1 ~]# razor nodes node1 facts
From http://localhost:8080/api/collections/nodes/node1:

    hardwareisa: x86_64
    macaddress: 00:0c:29:04:a1:a3
    architecture: x86_64
    hardwaremodel: x86_64
    processor0: Intel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz
    processorcount: 1
    interfaces: ens32,ens33,lo
    macaddress_ens32: 00:0c:29:04:a1:a3
    ipaddress_ens33: 192.168.1.110
    macaddress_ens33: 00:0c:29:04:a1:ad
    netmask_ens33: 255.255.255.0
    ipaddress_lo: 127.0.0.1
    netmask_lo: 255.0.0.0
    memorysize_mb: 365.05
    memoryfree_mb: 262.44
    facterversion: 2.0.1
    ipaddress: 192.168.1.110
    virtual: vmware
    is_virtual: true
    uniqueid: 007f0100
    blockdevice_sda_size: 19327352832
    blockdevice_sda_vendor: VMware
    blockdevice_sda_model: Virtual disk
    blockdevice_sr0_size: 1073741312
    blockdevice_sr0_vendor: NECVMWar
    blockdevice_sr0_model: VMware IDE CDR10
    blockdevices: sda,sr0
```

Figure 8: Razor Node Facts

```
img http://puppetagent1.cse.com:8080/svc/repo/CentOS6.5-Repo//images/product_img
...
Retrieving http://puppetagent1.cse.com:8080/svc/repo/CentOS6.5-Repo//images/i
nstaill.img...

Running anaconda 13.21.215, the CentOS system installer - please wait.
Examining storage devices
In progress

Creating device /dev/sdal
Creating ext4 filesystem on /dev/sdal
In progress


Creating device /dev/sda2
Creating physical volume (LVM) on /dev/sda2
In progress

Creating device /dev/mapper/vg_hostid
In progress

Creating device /dev/mapper/vg_hostid-lv_root
In progress
```

Figure 9: Automatic OS installation process

## Wrapping up

Razor provides real-time inventory data for every hardware node. Its auto-discovered ability eliminates the inefficient, error-prone manual process. Razor effectively uses IT defined policy to specify the desired state of each hardware node and its operating system. It automatically tracks provisioning progress toward this state, and can even decide when to re-provision. This gives you full control over a node's boot sequence and a complete log of its lifecycle. With RESTful open APIs, Razor gives you full programmatic control of the rules and models that govern operating system image selection and hardware provisioning.  **END**

## References

- [1] Razor: <https://github.com/puppetlabs/razor-server>
- [2] Razor microkernel: <https://github.com/puppetlabs/razor-el-mk>
- [3] Next Gen Provisioning: <http://puppetlabs.com/solutions/next-generation-provisioning>

## By: Ajeet S Raina

The author is a senior systems engineer currently working in R&D at Dell. He likes blogging on Big Data, OpenStack, virtualisation and Linux. You can follow his work through <http://collabnix.com>.



# SECURE YOUR LINUX BOX WITH EFFECTIVE USER AND GROUP MANAGEMENT

Security is a paramount concern today with identity theft becoming an everyday occurrence. Linux security is built around privileges given to users and groups. This article introduces the reader to the basics of user and group management in Linux.



Linux is basically a kernel on which various distributions of operating systems have been developed, namely Red Hat, Fedora, Ubuntu, SUSE, etc. Since its inception, Linux has become the most dominant technology used to manage servers worldwide. In recent years, its popularity has extended to normal users also, rather than just being the administrators' preferred choice. No doubt, one of the compelling attractions of Linux is that it is open source but another reason for choosing Linux for middle to high-end machines is that Linux is safe and secure.

A common myth is that Linux is built up on only plain text files, and that it is 100 per cent virus-free. However, this is not the truth. Hackers and crackers have always tried to inject threats into the Linux environment, and initially, they were successful to a certain extent. The reason for Linux's security is that all its processes run strictly under the privileges allocated to various users by the systems administrator. Privileges, if implemented fairly well, make an unbreakable security layer around the Linux engine and prevent it from being attacked. In Linux, users and group members access a file system's contents (files and directories) based on the privileges assigned to them in the

form of permissions. This is further enhanced by Access Control Lists, Sticky Bits and Security Enhanced Linux (SELinux).

## Users and groups

A **user** is a person who is authorised to use the Linux interface under privileges issued by the systems administrator. Every person gets a username and password as credentials, which are to be passed during the login process for receiving an interactive shell. When we add a new user, several details are automatically configured by Linux. However, these can be customised as we will see later. There are five kinds of users:

- **System user** – works for system applications and has a non-interactive shell. These users need not log in; rather, they become active as their corresponding system services start.
- **Super user** – is the one who has full control over the Linux file system and is empowered with unlimited permissions. Root is the default super user.
- **Owner user** – is the one who is the creator or owner of the content and uses the allotted permissions.

- *Group user* - is one member of a group, all the members of which get the same permissions for some particular content.
- *Other user* – is the one who is not an owner of content. *Write* and *Execute* permissions are not given to this user, unless really required.

A **group** is a collection of users to whom common permissions are allocated for any content. There is only one user owner for any content but, when we need to allow multiple users to access or modify the content by working as shared owners of that content, group ownership comes into the picture.

There are two types of groups:

- Primary group is the default group of a user, which is mandatory and grants group ownership. It is mapped in */etc/passwd* file with the corresponding user.
- The secondary group allows multiple users to become members in order to provide them content-accessing permissions. Members of this group do not get ownership of content.



**Tips: 1.** Adding and managing users seems to be very simple in graphical mode, but for administrators it is recommended to use commands on the terminal.

**2.** One user can have only one primary group, which is called his private group. However, users can enjoy secondary membership of multiple groups. This means that a user can be a member of multiple groups but only through secondary group membership.

## Creating users and groups with default configurations

Let us add two new users named *u1* and *u2* with default configurations and issue a password to *u1* only.

```
#useradd <username>
#passwd <username>
```

For Example:

```
#useradd u1
#passwd u1
#useradd u2
```

By default, a new group is created with the same name as the username and a */bin/bash* shell is assigned. Also, the user gets a directory with this name inside */home* and receives some predefined files from the */etc/skel* directory for starting work.

A new group can be added as follows:

```
#groupadd <groupname>
```

For example:

```
#groupadd mygroup
```

## Information and configuration parameters of users and groups

Information about existing users is present in */etc/passwd* and has the following format:

```
<username>:<auth_check>:<userid>:<groupid>:<Comments>:<Homed  
irectory>:<Shell>
```

The various fields separated by ‘:’ in this file are:

- *username* is a unique name issued to the person who is to use the Linux interface.
- The *auth\_check* field only denotes whether this user has a password or not. The actual password is stored in */etc/shadow* file.
- The *userid* is a unique number issued to identify a user. 1 - 499 are reserved for system applications and 499 onwards are issued to new users managed by the administrator. With default configurations, this ID is issued in an incremental manner.
- The *groupid* is a unique number issued to identify a group of users. 1- 499 are reserved for system applications and 499 onwards are issued to new groups managed by the administrator. With default configurations, this ID is issued in an incremental manner through the */etc/group* file.
- The *Comments* field is optional and contains information about the user.
- The *Home directory* is the default path where users can manage their documents and keep them safe from other users.
- *Shell* is an environment provided as an interface to the user. *Shell* decides which tools, utilities and facilities are to be provided. Examples are ksh, bash, nologin, etc.

Password related information of the user is automatically managed in the */etc/shadow* file. Here, passwords are saved in encrypted form.



**Tips:** Log in as the root user (the superuser with unlimited privileges) to access and manage these files.

Let us have a look at the */etc/passwd* file by using the *tail* command.

You can see that system user ‘*tcpdump*’ has IDs less than 500 while *u1* and *u2* have IDs that are greater than 500.

Let us have a look at the */etc/shadow* file by using the *tail* command.

Here, *u1* has an encrypted password while *tcpdump* and *u2* have null passwords.

```
[root@localhost Desktop]# tail -4 /etc/shadow
sshd:!!:16388:::::::
tcpdump:!!:16388:::::::
u1:$1$j0qyg2wK$EdiKHYtaGH7awnQVUR5Cg.:16388:0:99999:7:::
u2:!!:16388:0:99999:7:::
[root@localhost Desktop]#
```

Figure 1: Output of the tail command

```
[root@localhost Desktop]# tail -4 /etc/group
slocate:x:21:
u1:x:500:
u2:x:501:
mygroup:x:502:
[root@localhost Desktop]# █
```

Figure 2: /etc/group



**Tip:** `/etc/passwd-` and `/etc/shadow-` are the backup files of `/etc/passwd` and `/etc/shadow`, respectively.

You can switch between user sessions by using the `su` command as `#su - <username>`

For example:

```
#su - u1
```

Home directories of users exist under `/home`. The same can be checked by listing the contents of `/home` directory using command `#ls /home`. Directories of `u1` and `u2` will be displayed.

Information about existing groups is present in the `/etc/group` file and has the following format:

```
<groupname>:< auth_check >:<groupid>:<secondarymemberslist>
```

Let us have a look at the `/etc/group` file (Figure 2).

Here, no group is allotted as secondary to any user but soon we will see how to allocate secondary groups.

Group passwords are stored in the `/etc/gshadow` file.

Now the question is: where do the rules of default configurations reside? This is important for customisation.

The following files answer this question:

1. `/etc/login.defs`: Provides control over password aging, minimum-maximum allowable user IDs/group IDs, home directory creation and other settings.
2. `/etc/default/useradd`: Provides control over the home directory path, shell to be provided, kit-providing directory and other settings.
3. `/etc/skel`: This is the default directory acting as a start-up kit, which provides some pre-created files for new users.

## Customising default configurations for a new user

It is always good to tweak the default configurations for several reasons. It provides better security since the default settings are known to almost every admin. It also creates a customised environment that suits your requirements. It even helps you to exhibit your advanced administration skills.

Let us customise the following default settings for any new user:

- Change the default home directory path.
- Change the permitted maximum `user/group` ID range.
- Change the location of the default start-up kit source.

To do this, edit `login.defs` and make the changes as shown:

```
# vim /etc/login.defs
```

Now, edit the `useradd` file:

```
# vim /etc/default/useradd
```

Next, create the `/ghar` and `/etc/startkit` directories with some files in `startkit`.

So, we have created a file named `Welcome`, which will be available in the home directory of all new users. Now, add a new user `u3` and check its home directory under `/ghar` instead of `/home`.

We can see that `u3` received the file named `Welcome` in the home directory automatically. Now, when we try to add more users, the effect of `Min` and `Max` `user/group` ids will not allow us to accomplish `user/group` creation. We cannot create more users since the maximum range of the UID has been reached, which can be verified in `/etc/passwd`.

Other alterations are also visible in `/etc/passwd` as per customisations.

## Dynamically customising default configurations for a new user

We can also make changes in configurations for individual users during new user creation. The following are the options for `useradd` and `usermod` commands:

```
-u -> uid-d -> home directory -g -> primary group
-s -> shell -G -> Secondary Group
```

Let us add a new user called 'myuser' with the user ID as '2014', the home directory as '/myhome', the primary group as 'u1', secondary group as 'u2' and shell as 'ksh'. Any of these options can be omitted, as required.

```
# useradd -u 2014 -d /myhome -g u1 -G u2 -s /sbin/ksh
myuser
```



**Tip:** Manual UID and GID do not adhere to minimum and maximum ID rules.

Let us verify the user information in `/etc/passwd`.

Also, look at the `/etc/group` file to see the effect of secondary group membership of 'myuser' with 'u2'.



**Tip:** After customising the primary group through the command, a new group with the name of the user is not created.

Similarly, a group can also be configured with a specific ID during creation:



```
[root@localhost etc]# useradd u3
[root@localhost etc]# ls /home
u1  u2
[root@localhost etc]# ls /ghar
u3
[root@localhost etc]# ls /ghar/u3
Welcome
[root@localhost etc]# █
```

Figure 3: Listing of the home directory

```
# groupadd -g 2222 testgroup
```

By default, every group should have a unique group ID, but a non-unique group ID can be forcibly shared among different group names using `-o` as follows:

```
# groupadd -o -g 2222 newgroup
```

We can verify the result of above commands in `/etc/group`.

## Altering existing users and groups

The `usermod` and `groupmod` commands (with options that are the same used in creation) are used to alter the information, while `userdel` and `groupdel` commands are used to delete existing users and groups, respectively.

To change the user ID and secondary group of 'myuser', type:

```
# usermod -u 3333 -G newgroup myuser
```

Now, verify `/etc/passwd` and `/etc/group` files.

To change the name of 'newgroup' to 'ngroup', type:

```
# groupmod -n ngroup newgroup
```

To modify 'gid' of 'ngroup' from '2222' to '4444', type:

```
# groupmod -g 4444 ngroup
```

To delete an existing group – 'testgroup', type:

```
# groupdel testgroup
```

To delete an existing user 'u2', type:

```
# userdel u2
```

To delete an existing user 'u2' along with its home directory, type:

```
# userdel -r u2 
```

### By: Shobhit Gupta

The author is an academician-cum-Developer-cum-Administrator having more than 5 years of IT experience with commitment to share knowledge through two books and several journal papers/articles published and providing quality service. He can be contacted at [shobhit.gupta.it@gmail.com](mailto:shobhit.gupta.it@gmail.com)

## OSFY Magazine Attractions During 2014-15

| MONTH          | THEME                               | FEATURED LIST                      | BUYERS' GUIDE                |
|----------------|-------------------------------------|------------------------------------|------------------------------|
| March 2014     | Network monitoring                  | Security                           | -----                        |
| April 2014     | Android Special                     | Anti Virus                         | Wifi Hotspot Devices         |
| May 2014       | Backup and Data Storage             | Certification                      | External Storage             |
| June 2014      | Open Source on Windows              | Mobile Apps                        | UTMs fo SMEs                 |
| July 2014      | Firewall and Network security       | Web Hosting Solutions Providers    | MFD Printers for SMEs        |
| August 2014    | Kernel Development                  | Big Data Solutions Providers       | SSDs for Servers             |
| September 2014 | Open Source for Start-ups           | Cloud                              | Android Devices              |
| October 2014   | Mobile App Development              | Training on Programming Languages  | Projectors                   |
| November 2014  | Cloud Special                       | Virtualisation Solutions Providers | Network Switches and Routers |
| December 2014  | Web Development                     | Leading Ecommerce Sites            | AV Conferencing              |
| January 2015   | Programming Languages               | IT Consultancy Service Providers   | Laser Printers for SMEs      |
| February 2015  | Top 10 of Everything on Open Source | Storage Solutions Providers        | Wireless Routers             |



# Automate Repetitive Tasks with SaltStack

SaltStack is an open source, scalable, flexible systems management software that can be used for data centre automation, cloud orchestration, server provisioning and configuration management.



**C**M or configuration management is common terminology in the software industry. For those of us who do not know what it is, it's the process of automating sets of tasks that are repetitive, e.g., building a machine to match the company or group standards. Some of the prominent tools that help you to do this are SaltStack, Ansible, Puppet, Chef, CFEngine, Bcfg2, Quattor, Rundeck and Vagrant. We will concentrate on SaltStack.

A DevOps engineer is a new role that has evolved in the software industry. One of the key skills a DevOps engineer is expected to have, among many others, is automation skills. Whoever has read the four-part series of articles on Puppet at <http://www.opensourceforu.com/tag/puppet/> must surely know about configuration management by now, at least from a Puppet perspective. But don't worry if you don't know too much; one can always learn.

## Why SaltStack?

There are a few reasons why we have chosen SaltStack as a configuration management tool. These are:

- Orchestration, predictability and automation
- The power available to DevOps
- It is open source (Apache licence), free, and has pretty wide community support in case of hiccups

- A full-fledged client/server model
- It is written entirely in Python
- Marks systems for easier management
- Uses the fast, powerful and light communication layer—ZeroMQ
- Uses the well established jinja2 platform for logic (uses a variety of languages to render the configuration files; yaml + jinja2 is the default template engine)
- Password-less root-based remote execution—runs predefined or custom commands on remote hosts
- It uses a push technology, where the onus is on the server to push configuration details to the client
- Completely encrypted communication and data transfer
- It is available on all UNIX/GNU/Linux platforms under the following pkg names:
  - salt-master
  - salt-minion

## Glossary

This is already available at the SaltStack glossary page; I am just re-pasting a few terms that we will see more often.

**SLS:** The SaLt State File Extension.

**Salt State:** Consists of identifiers (resources to operate on), a state module (functionality to manage a type of resource) and function declarations (definition of a desired state of a resource), and sometimes function parameters (specific details for the state function).

**Master:** A central Salt daemon, from which commands can be issued to listening minions.

**Minion:** A server running a Salt minion daemon, which can listen to commands from a master and perform the requested tasks. Generally, minions are servers that are to be controlled using Salt.

**Masterless:** A minion that does not require a Salt master to operate. All configuration is local.

**Top file:** Determines which SLS files should be applied to various systems, and organises those groups of systems into environments.

**High state:** The collection of states to be applied to a system.

**High data:** Data structure in a SLS file that represents a

## How SaltStack compares with Puppet

|                                      | Salt  | Puppet  |
|--------------------------------------|---|---|
| File extensions                      | .sls  | .pp   |
| Base language                        | Jinja2-based YAML   | Ruby-based DSL  |
| Root structure                       | file roots (has options of Dev, Stage or Prod)  | module path   |
| Entry point                          | top.sls or init.sls   | site.pp or init.pp  |
| Client                               | Minion  | Client  |
| Augeas module for file modifications | Not present   | Present   |
| Dashboard                            | No visual interface   | Available in enterprise edition   |
| Templating structure                 | Not that strong   | Powerful  |
| Dynamic Global Variables Declaration | Orderly but reduces flexibility; resides in separate location and cannot be declared on the fly. Variables are called pillars | More intuitive and flexible   |
| Orchestration                        | ZeroMQ  | Mcollective with no default functionality   |
| Remote execution                     | Out-of-the-box features, supports dynamic querying and scalable orchestration   | There is no option, as of now   |
| Iterations                           | Loops within code are very flexible   | Not that flexible; in fact, it is very tough to code loops for re-use of variables or call more than one value at the same time |

set of state declarations.

**Job:** Tasks to be performed by Salt command execution.

**Pillar:** A simple key-value store for user-defined data to be made available to a minion. Often used to store and distribute sensitive data to minions.

**Grain:** A key-value pair which contains a fact about a system, such as its hostname or network addresses.

**Jinja:** A templating language framework for Python, which allows variables and simple logic to be dynamically inserted into static text files when they are rendered. Inspired by Django's templating language.

**Salt key:** Salt master manages which machines are allowed or not allowed to communicate with it.

**Salt SSH:** A configuration management and remote orchestration system, which does not require that any software besides SSH be installed on systems to be controlled.

**SLS module:** Contains a set of state declarations.

**State declaration:** A data structure that contains a unique ID and describes one or more states of a system, such as ensuring that a package is installed or a user is defined.

**State module:** A module that contains a set of state functions.

**State run:** The application of a set of states on a set of systems.

**Target:** Minion(s) to which a given Salt command will apply.

## Setup

We can install SaltStack by three methods—Bootstrap, Yum or APT. We are only covering installs for the most common server platforms like Red Hat, SUSE and Debian, though the SaltStack official documentation includes Solaris, ArchLinux, Ubuntu and even Windows.

**Bootstrap:** This is the easiest way and it takes care of

dependency packages as well on any platform.

Download the required script:

```
wget --no-check-certificate -O install_salt.sh http://bootstrap.saltstack.org
```

Install the master:

```
sh install_salt.sh -M -N
```

Install the client:

```
sh install_salt.sh -A [Master's IP or DNS Record]
```

**Yum:** Yum is platform-independent and is best when we use Red Hat or SUSE. It will install all the required dependencies.

Install the master:

```
yum install salt-master
```

Now, install the client:

```
yum install salt-minion
```

For SUSE, replacing Yum with Zypper should also work.

```
zypper addrepo http://download.opensuse.org/repositories/devel:languages:python/SLE_11_SP3/devel:languages:zypper refresh
zypper install salt salt-minion salt-master
```



**APT:** APT is best for Debian platforms and it takes care of all the dependencies as well.

Add the following to */etc/apt/sources.list*

```
deb http://debian.saltstack.com/debian squeeze-saltstack main
deb http://backports.debian.org/debian-backports squeeze-
backports main contrib non-free
```

And import the required key before installing the following:

```
wget -q -O- "http://debian.saltstack.com/debian-salt-team-
joehealy.gpg.key" | apt-key add -
apt-get install salt-master
apt-get install salt-minion
```

**TAR:** This is the least preferred way because of the dependencies involved, but is do-able in the following way:

```
wget https://pypi.python.org/packages/source/s/salt/salt-
2014.1.10.tar.gz --no-check-certificate
tar -zxvf salt-2014.1.10.tar.gz
cd salt-2014.1.10
python26 setup.py install
wget https://bootstrap.pypa.io/get-pip.py --no-check-
certificate (for pip command for python26)
```

**Dependencies for TAR:** Using PIP or Source Tar, the following packages have to be installed for SaltStack and Python to go hand-in-hand:

```
yum install python26-jinja2.noarch python26-PyYAML python26-
zmq.x86_64 python26-m2crypto.x86_64 python26-msgpack.x86_64
python26-crypto.x86_64 python26-PyYAML.x86_64
/usr/bin/pip2.6 install msgpack-pure
salt-master --versions-report
```

```
Salt: 2014.1.10
Python: 2.6.8 (unknown, Nov 7 2012, 14:47:45)
Jinja2: 2.5.5
M2Crypto: 0.21.1
msgpack-python: 0.1.12
msgpack-pure: 0.1.3
pycrypto: 2.3
PyYAML: 3.08
PyZMQ: 2.1.9
ZMQ: 2.1.9
```

## Configuration

### Master:

- If it is a tar-based install, create the folder */etc/salt* and copy the file *conf/master* from the *Salt install* folder to */etc/salt/master*. In other cases, the file should be there by default.
- By default, it will take the path */srv/salt* and */srv/pillar* as the base folders, if not mentioned. But still, we can go and

uncomment the following lines in */etc/salt/master*

- `file_roots:`
- `base:`
- `- /srv/salt`

- Run the command *salt-master* with a *-d* if you want to run it as a background daemon.

### Client:

- If it is a tar-based install, create the folder */etc/salt* and copy the file *conf/minion* from the *salt install* folder to */etc/salt/minion*. In other cases, the file should be there, by default.
- Add *<ip-address> salt salt.master.com* to */etc/hosts*. Only and only if DNS doesn't have entry for the Salt master, do we need to add an entry here that resolves entries to the word 'salt'.
- Search for *"#master:"* in the file */etc/salt/minion*; uncomment it before adding the name of the Salt master after the colon and saving the file.
- Run the command *salt-minion* with a *-d* if you want to run it as a background daemon. This will send an authorisation request to the master.

### Activation:


- From the server, list the machines waiting for authorisation using:
  - `salt-key -L`
- To authorise a node, sign the respective node
  - `salt-key -a 'node-name'`
  - `salt-key -a` (accepts all nodes)

### Testing:

- Run a test from the master to the client.
  - `salt '*' test.ping`
  - which gives output as...
    - `Node_name:`
    - `True`

### Uninstalling:

```
yum remove salt salt-master salt-minion (RedHat or SuSE)
zypper remove salt salt-master salt-minion (SuSE)
apt-get autoremove salt salt-master salt-minion (Debian, Ubuntu)
```

If we haven't removed the untarred *salt-tar* folder, then a *make clean* must clean things up. 

## Resources

- [1] <http://salt.readthedocs.org/en/v0.9.2/topics/download.html>
- [2] <http://docs.saltstack.com/en/latest/>
- [3] <http://www.linuxjournal.com/content/getting-started-salt-stack-other-configuration-management-system-built-python>

## By: Bejoy Abraham Mathews and Davana Satish Gongadi

Bejoy Abraham Mathews works with Servista Inc. He has earlier authored articles on Puppet and is RHCE, the cloud and Big Data certified. He can be contacted at [bejoy.abraham@gmail.com](mailto:bejoy.abraham@gmail.com)

Davana Satish Gongadi works with Quadrant 4 System Corporation. He is SCJP certified and is an expert in Java Application Server environments on GNU/Linux. He can be contacted at [gd\\_satish@yahoo.com](mailto:gd_satish@yahoo.com).

# Playing Around with Graphs in Maxima

In the final and 24<sup>th</sup> article in this series, the reader gets to play around with graphs using previously gained knowledge about the graphs package of Maxima.

In the previous article in this series, we got familiar with simple graphs, and how the graphs package of Maxima allows us to create and visualise them. Building on that knowledge, in this article, we are going to play around with graphs and their properties, using the functions provided by Maxima's graphs package.

## Graph modifications

We have already created various graphs with the *create\_graph()* and *make\_graph()* functions of the graphs package of Maxima. What if we wanted to modify some existing graphs, say by adding or removing some edges or vertices? For such operations, Maxima provides the following functions:

- *add\_edge(<edge>, <g>)* - Adds <edge> into the graph <g>
- *add\_edges(<edge\_list>, <g>)* - Adds edges specified by <edge\_list> into the graph <g>
- *add\_vertex(<vertex>, <g>)* - Adds <vertex> into the graph <g>
- *add\_vertices(<vertex\_list>, <g>)* - Adds vertices specified by <vertex\_list> into the graph <g>
- *connect\_vertices(<u\_list>, <v\_list>, <g>)* - Connects all vertices from <u\_list> to all vertices in <v\_list> in the graph <g>
- *contract\_edge(<edge>, <g>)* - Merges the vertices of the <edge> and the edges incident on those vertices, in the graph <g>
- *remove\_edge(<edge>, <g>)* - Removes the <edge> from the graph <g>
- *remove\_vertex(<vertex>, <g>)* - Removes the <vertex> and the associated edges from the graph <g>

Some of the above functions are demonstrated below:

```
$ maxima -q
(%i1) load(graphs)$ /* Loading the graphs package */
...
0 errors, 0 warnings
(%i2) g: create_graph(4, [[0, 1], [0, 2]]);
(%o2) GRAPH(4 vertices, 2 edges)
(%i3) print_graph(g)$
```

Graph on 4 vertices with 2 edges.

```
Adjacencies:
3 :
2 : 0
1 : 0
0 : 2 1
(%i4) add_edge([1, 2], g)$
(%i5) print_graph(g)$
```

Graph on 4 vertices with 3 edges.

```
Adjacencies:
3 :
2 : 1 0
1 : 2 0
0 : 2 1
(%i6) contract_edge([0, 1], g)$
(%i7) print_graph(g)$
```

Graph on 3 vertices with 1 edges.

Adjacencies:

```
3 :
2 : 0
0 : 2
```

In the above examples, if we do not intend to modify the original graph, we can make a copy of it using *copy\_graph()*, and then operate on the copy, as follows:

```
(%i8) h: copy_graph(g);
(%o8)          GRAPH(3 vertices, 1 edges)
(%i9) add_vertex(1, h)$
(%i10) print_graph(h)$
```

Graph on 4 vertices with 1 edges.

Adjacencies:

```
1 :
0 : 2
2 : 0
3 :
```

```
(%i11) print_graph(g)$ /* Notice g is unmodified */
```

Graph on 3 vertices with 1 edges.

Adjacencies:

```
3 :
2 : 0
0 : 2
```

```
(%i12) quit();
```

## Advanced graph creations

New graphs can also be created based on existing graphs and their properties by various interesting operations. A few of them are listed below:

- *underlying\_graph(<dg>)* - Returns the underlying graph of the directed graph <dg>
- *complement\_graph(<g>)* - Returns the complement graph of graph <g>
- *line\_graph(<g>)* - Returns a graph that represents the adjacencies between the edges of graph <g>
- *graph\_union(<g1>, <g2>)* - Returns a graph with edges and vertices of both graphs <g1> and <g2>
- *graph\_product(<g1>, <g2>)* - Returns the Cartesian product of graphs <g1> and <g2>

Here are some examples to demonstrate the simpler functions:

```
$ maxima -q
(%i1) load(graphs)$
...
0 errors, 0 warnings
(%i2) g: create_graph(4, [[0, 1], [0, 2], [0, 3]], directed =
true);
(%o2)          DIGRAPH(4 vertices, 3 arcs)
(%i3) print_graph(g)$
```

Digraph on 4 vertices with 3 arcs.

Adjacencies:

```
3 :
2 :
1 :
0 : 3 2 1
```

```
(%i4) h: underlying_graph(g);
(%o4)          GRAPH(4 vertices, 3 edges)
(%i5) print_graph(h)$
```

Graph on 4 vertices with 3 edges.

Adjacencies:

```
0 : 1 2 3
1 : 0
2 : 0
3 : 0
```

```
(%i6) print_graph(complement_graph(h))$
```

Graph on 4 vertices with 3 edges.

Adjacencies:

```
3 : 2 1
2 : 3 1
1 : 3 2
0 :
```

```
(%i7) print_graph(graph_union(h, complement_graph(h)))$
```

Graph on 8 vertices with 6 edges.

Adjacencies:

```
4 :
5 : 6 7
6 : 5 7
7 : 5 6
3 : 0
2 : 0
1 : 0
0 : 3 2 1
```

```
(%i8) quit();
```

## Basic graph properties

*graph\_order(<g>)*, *vertices(<g>)* returns the number of vertices and the list of vertices, respectively, in the graph <g>. *graph\_size(<g>)*, *edges(<g>)* returns the number of edges and the list of edges, respectively, in the graph <g>.

A graph is a collection of vertices and edges. Hence, most of its properties are centred around them. The following are graph related predicates provided by the graphs package of Maxima:

- *is\_graph(<g>)* - returns 'true' if <g> is a graph, and 'false' otherwise
- *is\_digraph(<g>)* - returns 'true' if <g> is a directed graph, and 'false' otherwise
- *is\_graph\_or\_digraph(<g>)* - returns 'true' if <g> is a graph or a directed graph, and 'false' otherwise
- *is\_connected(<g>)* - returns 'true' if graph <g> is



connected, and 'false' otherwise

- *is\_planar(<g>)* - returns 'true' if graph <g> can be placed on a plane without its edges crossing each other, and 'false' otherwise
- *is\_tree(<g>)* - returns 'true' if graph <g> has no simple cycles, and 'false' otherwise
- *is\_biconnected(<g>)* - returns 'true' if graph <g> will remain connected even after removal of any one of its vertices and the edges incident on that vertex, and 'false' otherwise
- *is\_bipartite(<g>)* - returns 'true' if graph <g> is bipartite, i.e., two-colourable, and false otherwise
- *is\_isomorphic(<g1>, <g2>)* - returns 'true' if graphs <g1> and <g2> have the same number of vertices and are connected in the same way, and 'false' otherwise. And, isomorphism (<g1>, <g2>) returns an isomorphism (that is a one-to-one onto mapping) between the graphs <g1> and <g2>, if it exists.
- *is\_edge\_in\_graph(<edge>, <g>)* - returns 'true' if <edge> is in graph <g>, and 'false' otherwise
- *is\_vertex\_in\_graph(<vertex>, <g>)* - returns 'true' if <vertex> is in graph <g>, and 'false' otherwise

The following example specifically demonstrates the isomorphism property, from the list above:

```
$ maxima -q
(%i1) load(graphs)$
...
0 errors, 0 warnings
(%i2) g1: create_graph(3, [[0, 1], [0, 2]]);
(%o2) GRAPH(3 vertices, 2 edges)
(%i3) g2: create_graph(3, [[1, 2], [0, 2]]);
(%o3) GRAPH(3 vertices, 2 edges)
(%i4) is_isomorphic(g1, g2);
(%o4) true
(%i5) isomorphism(g1, g2);
(%o5) [2 -> 0, 1 -> 1, 0 -> 2]
(%i6) quit();
```

## Graph neighbourhoods

A lot of the properties of graphs are linked to vertex and edge neighbourhoods, also referred to as adjacencies.

For example, a graph itself could be represented by an adjacency list or matrix, which specifies the vertices adjacent to the various vertices in the graph. *adjacency\_matrix(<g>)* returns the adjacency matrix of the graph <g>.

The number of edges incident on a vertex is called the valency or degree of the vertex, and could be obtained using *vertex\_degree(<v>, <g>)*. *degree\_sequence(<g>)* returns the list of degrees of all the vertices of the graph <g>. In case of a directed graph, the degrees could be segregated as in-degree and out-degree, as per the edges incident into and out of the vertex, respectively. *vertex\_in\_degree(<v>, <dg>)* and *vertex\_out\_degree(<v>, <dg>)*, respectively, return the in-degree and out-degree for the vertex <v> of the directed graph <dg>.

*neighbors(<v>, <g>)*, *in\_neighbors(<v>, <dg>)* and *out\_neighbors(<v>, <dg>)* return the list of adjacent vertices,

adjacent in-vertices and the adjacent out-vertices, respectively, of the vertex <v> in the corresponding graphs.

*average\_degree(<g>)* computes the average of the degrees of all the vertices of the graph <g>. *max\_degree(<g>)* finds the maximal degree of vertices of the graph <g>, and returns one such vertex alongwith. *min\_degree(<g>)* finds the minimal degree of vertices of the graph <g>, and returns one such vertex alongwith.

Here follows a neighbourhood related demonstration:

```
$ maxima -q
(%i1) load(graphs)$
...
0 errors, 0 warnings
(%i2) g: create_graph(4, [[0, 1], [0, 2], [0, 3], [1, 2]]);
(%o2) GRAPH(4 vertices, 4 edges)
(%i3) string(adjacency_matrix(g)); /* string for output in single line */
(%o3) matrix([0,0,0,1],[0,0,1,1],[0,1,0,1],[1,1,1,0])
(%i4) degree_sequence(g);
(%o4) [1, 2, 2, 3]
(%i5) average_degree(g);
(%o5) 2
(%i6) neighbors(0, g);
(%o6) [3, 2, 1]
(%i7) quit();
```

## Graph connectivity

A graph is ultimately about connections, and hence lots of graph properties are centred around connectivity.

*vertex\_connectivity(<g>)* returns the minimum number of vertices that need to be removed from the graph <g> to make the graph <g> disconnected. Similarly, *edge\_connectivity(<g>)* returns the minimum number of edges that need to be removed from the graph <g> to make the graph <g> disconnected.

*vertex\_distance(<u>, <v>, <g>)* returns the length of the shortest path between the vertices <u> and <v> in the graph <g>. The actual path could be obtained using *shortest\_path(<u>, <v>, <g>)*.

*girth(<g>)* returns the length of the shortest cycle in graph <g>.

*vertex\_eccentricity(<v>, <g>)* returns the maximum of the vertex distances of vertex <v> with any other vertex in the connected graph <g>.

*diameter(<g>)* returns the maximum of the vertex eccentricities of all the vertices in the connected graph <g>.

*radius(<g>)* returns the minimum of the vertex eccentricities of all the vertices in the connected graph <g>.

*graph\_center(<g>)* returns the list of vertices that have eccentricities equal to the radius of the connected graph <g>.

*graph\_periphery(<g>)* is the list of vertices that have eccentricities equal to the diameter of the connected graph.

A minimal connectivity related demonstration for the graph shown in Figure 1 follows:

```
$ maxima -q
(%i1) load(graphs)$
```

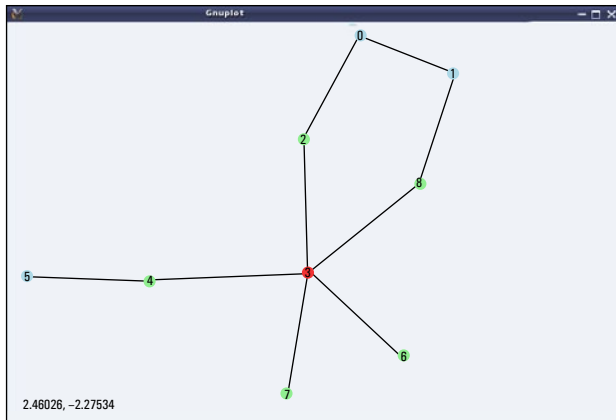


Figure 1: Graph connectivities

```
...
0 errors, 0 warnings
(%i2) g: create_graph(9, [[0, 1], [0, 2], [1, 8], [8, 3], [2, 3],
[3, 4], [4, 5], [3, 6], [3, 7]]);
(%o2) GRAPH(9 vertices, 9 edges)
(%i3) vertex_connectivity(g);
(%o3) 1
(%i4) edge_connectivity(g);
(%o4) 1
(%i5) shortest_path(0, 7, g);
(%o5) [0, 2, 3, 7]
(%i6) vertex_distance(0, 7, g);
(%o6) 3
(%i7) girth(g);
(%o7) 5
(%i8) diameter(g);
(%o8) 4
(%i9) radius(g);
(%o9) 2
(%i10) graph_center(g);
(%o10) [3]
(%i11) graph_periphery(g);
(%o11) [5, 1, 0]
```

Vertex 3 is the only centre of the graph, and 0, 1 and 5 are the peripheral vertices of the graph.

## Graph colouring

Graph colouring has been a fascinating topic in graph theory, right since its inception. It is all about colouring the vertices or edges of a graph in such a way that no adjacent elements (vertex or edge) have the same colour.

The smallest number of colours needed to colour the vertices of a graph, such that no two adjacent vertices have the same colour, is called the chromatic number of the graph. *chromatic\_number(<g>)* computes the same. *vertex\_coloring(<g>)* returns a list representing the colouring of the vertices of <g>, along with the chromatic number.

The smallest number of colours needed to colour the




Figure 2: Graph colouring


edges of a graph, such that no two adjacent edges have the same colour, is called the chromatic index of the graph. *chromatic\_index(<g>)* computes the same. *edge\_coloring(<g>)* returns a list representing the colouring of the edges of <g>, along with the chromatic index.

The following demonstration continues colouring the graph from the above demonstration:

```
(%i12) chromatic_number(g);
(%o12) 3
(%i13) vc: vertex_coloring(g);
(%o13) [3, [[0, 3], [1, 1], [2, 2], [3, 1], [4, 2], [5, 1], [6, 2], [7, 2], [8, 2]]]
(%i14) chromatic_index(g);
(%o14) 5
(%i15) ec: edge_coloring(g);
(%o15) [5, [[[0, 1], 1], [[0, 2], 2], [[1, 8], 2], [[3, 8], 5], [[2, 3], 1], [[3, 4], 2], [[4, 5], 1], [[3, 6], 3], [[3, 7], 4]]]
(%i16) draw_graph(g, vertex_coloring=vc, edge_coloring=ec, vertex_size=5, edge_width=3, show_id=true)$
(%i17) quit();
```

Figure 2 shows the coloured version of the graph, as obtained by %i16.

## Bon voyage

With this article, we have completed a two-year long mathematical odyssey through open source, starting from mathematics in Shell, covering Bench Calculator and Octave, and concluding with Maxima. I take this opportunity to thank my readers and wish them bon voyage with whatever they have gained through our interactions. However, this is not the end—get set for our next journey. **END** 

### By: Anil Kumar Pugalia

The author is a gold medallist from NIT Warangal and IISc Bengaluru. Mathematics and knowledge-sharing are two of his many passions. Learn more about him at <http://sysplay.in>. He can be reached at email@sarika-pugs.com.

# Speed Up your Research Paper with **zotero**



Zotero is an amazing open source tool that helps writers to manage the various references used in their articles and presentations. Essentially, it helps them to add these references to their articles with minimal drudgery but with clinical precision.

**P**ose this question to any research scholar, “What is the most laborious task while preparing a research paper or thesis?” More often than not, the answer will be *reference management*. It is true that managing references when writing research articles requires clinical precision. To beef up the paper, the references need to be handled in various formats for different journals, conferences, publishers, etc.

Managing references involves two major phases: the first is to organise the vast collection of bibliographic items gathered over a period of time, as a good library; the second is to search this library and cite them in manuscripts that you prepare. It wasn't long ago that all these things were carried out with very little intervention from software. Then came software that could assist both in managing the reference collection as well as in incorporating citations in the required format within the manuscripts. These software are called reference managers and an array of them, both proprietary and open source, are available now. This article highlights Zotero

—the leading open source reference manager.

## **Reference managers: Whom do they aid?**

Though reference managers are widely used by research scholars, this practice is not restricted to researchers and scientists alone. Reference managers can be used by any one preparing a project report, school or college assignment, or someone writing a thesis, a book or articles for reputed magazines. In general, reference managers are handy for all those whose writing involves citing the original sources of the article or book's contents.

A point to be noted is that publishing a paper or article by using already published content, without adding citations to the source, is known as plagiarism, which would lead to copyright infringement. In the research community, plagiarism is considered to be a serious issue. A reference manager makes your life simpler by adding and removing references swiftly.



## Why reference managers?

The reasons for using reference managers are listed below:

- Insertion and removal of citations becomes simple. In case of numbered citations, i.e., in the places where references are made to other material, numbers are inserted in brackets like [1],[2], etc. While revising the article, if you need to remove or add a citation, then all the following numbers need to be modified. You will understand the painstaking nature of this task, when the article being written involves a large number of references. If you are using a reference manager, then the numbers are incremented or decremented automatically.
- Another important reason for using a reference manager is to handle the reference list format properly. There are various format types like the APA style, MLA style, Chicago Manual of Style, the Harvard style, etc. Generally, different publishers follow different styles, and they mandate that the content writers adhere to these formats strictly. Hence, if you have to change from one format to another, it would consume hours of intensive and focused attention. Whereas, if you use reference managers, switching from one format to another can be done with a single click.
- Reference managers come handy in organising the collection of reference material gathered from various sources in a single window. Most of the reference managers have a search feature, which facilitates the search through the collections in the library.
- Reference managers enable collaboration with people of similar interests and also help in sharing resources. They facilitate syncing between multiple computers that an author may use.

## The choices

A quick Google search on reference managers will reveal that there are umpteen choices. Complete coverage of all reference managers is beyond the scope of this article. Figure 1 illustrates the prominent choices in the world of reference managers.

While the spectrum of reference managers provides users with options, it also makes choosing one a difficult task. Most reference managers provide the basic functionality of keeping the references and adding the citations. Apart from this, each reference manager has its own unique feature, which could be mobile support, search, collaboration, etc. So users have to make an informed decision while choosing a reference manager.

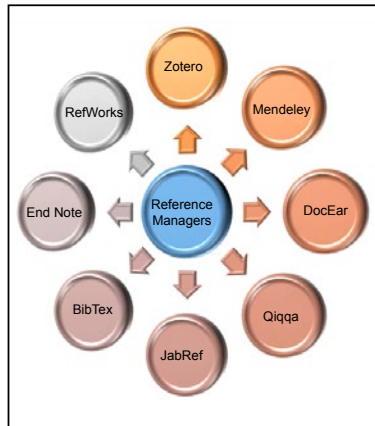


Figure 1: Reference manager choices




Figure 2: Factors to consider

## Factors to consider

Indeed, there are many factors that need to be considered.

Figure 2 illustrates various factors that affect the choice of a reference manager.

- **Cost:** As with any other decision-making process, cost plays an important role in selecting a reference manager if you plan to choose proprietary tools like EndNote, RefWorks, etc. Some of the reference managers have a hybrid plan. They provide some features for free and premium features need to be paid for. Mendeley, which was purchased by Elsevier in 2013, and Qiqa may be categorised under the hybrid category. If you opt for a free and open source product like Zotero, the cost factor becomes insignificant.
- **OS support:** If you happen to use multiple operating systems across various devices, then you have to check whether the reference manager supports all the OSs that you use. For example, EndNote supports Windows and Mac. If you are a Linux user, then this might become a bottleneck. Reference managers like Zotero, Docear, Mendeley, etc, support Windows, Linux and Mac.
- **Mobile phone support:** Smartphones have become a great productivity tool. If you want to work with your reference managers from anywhere, then mobile phone support needs to be considered. Reference managers like Qiqa provide an Android version. Mendeley versions are available for the iPad, iPhone, etc. For Zotero, there are applications like Scanner for Zotero, Zandy, etc, which are available in the Google Play store.
- **Word processor support:** Your manuscript will be created using the word processing tools of software suites like MS Office, Libre Office, etc. The reference manager's integration with word processing tools is a significant factor to consider. For example, Zotero supports various word processors like MS Word, Libre Office, Open Office, Google Docs, etc. If you are a LaTeX user, then reference managers like JabRef are handy.
- **Number of reference styles:** There are thousands of referencing styles. If you need to prepare manuscripts that

span a wider spectrum of styles, it is better to choose a reference manager that supports the maximum number of reference styles.

- **Search feature:** Once you have built a vast collection of references, searching across those collections becomes the most time-saving option.

- **PDF annotations:**

Reference managers like Mendeley have built-in PDF annotation features. Zotero has the facility to read the annotations made by other tools. Annotations are handy to make your comments on the PDF, which can be utilised later.

Apart from the above specified list, one must consider other features like browser support, import/export features, collaboration features, etc.

## Why Zotero?

From the list of selection factors specified above, it can be inferred that Zotero has many positive points in its favour like being cross-platform, open source, offering smartphone support, collaboration features, etc. Zotero is an open source reference manager that is popularly called a personal research assistant by its users. The initial release of Zotero was in the year 2006 by the Centre for History and New Media at George Mason University, Virginia, USA.

The primary reason for choosing Zotero is its simple and intuitive workflow. It provides five major features under a single umbrella, as shown in Figure 3.

## How to use Zotero

Now let us illustrate how to use Zotero with a series of steps. Zotero comes in two basic forms. One is the Zotero standalone version and the other is the Zotero Firefox add-on. The first version was launched as the Firefox add-on. The standalone version, which uses XULRunner, was launched in 2011. You can use either the Firefox add-on or the standalone




Figure 3: Why Zotero?

desktop version. They can be downloaded from the Zotero website <https://www.zotero.org/download/>

## Zotero standalone

The Zotero standalone looks like the screenshot shown in Figure 4. It lists all the reference items that you have gathered, under 'My Library' (marked as 1). We can create folders to group the items. The middle section (marked as 2) of the screen displays all the items under the selected folder. The last section (3) shows the fields under the selected item. It allows one to edit the fields as well. Tags and notes can also be added from the respective tabs.

## Zotero Firefox add-on

The Zotero Firefox add-on is used for the same purposes as the standalone version with the only difference being that the add-on version requires you to keep Firefox open whenever you want to use Zotero. The Firefox add-on also permits the items in the library to be modified. If you use both the versions in your system, then modifications from any one of them will reflect in the other as well.

Apart from the Firefox add-on, there are connectors available for many other browsers like Google Chrome, Opera, Safari, etc. If you prefer a browser other than Firefox, then a combination of this connector and standalone version is handy.

## Adding an item to Zotero

Adding an item to the Zotero library has been made very simple by the automatic content-sensing feature of Zotero. After installing Zotero standalone or the Firefox plugin, if you visit a publisher's site or any indexing site like Google Scholar, you will notice an icon in the address bar. By merely clicking the icon, the item is added to the library, as shown in Figure 5.

Any Web page can be added as an item to the Zotero library, by right-clicking and selecting *Save Zotero Snapshot* from the *Current Page*.

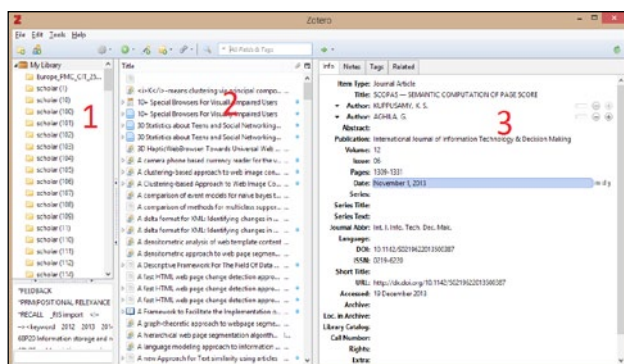


Figure 4: Screenshot of the Zotero standalone version

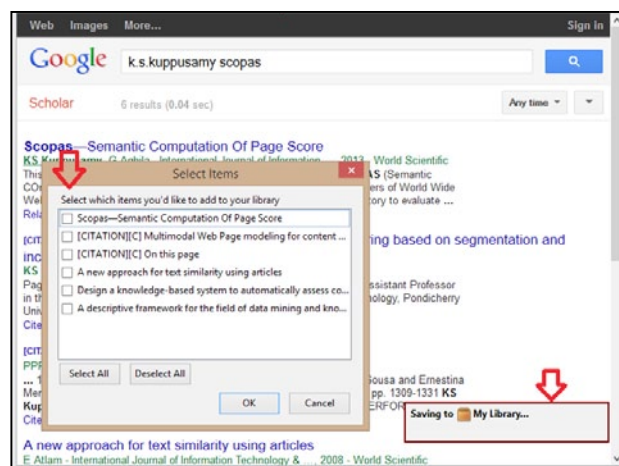


Figure 5: Adding an item to Zotero

## Word processor integration

To cite an item from your library, in the manuscript, the word processor plugin needs to be added. Word processor plugins are available for MS Word, Libre Office, etc. After installing the plugin, a new toolbar is added under the *Add-INS* menu. To add a citation, click on the first icon, *Zotero – Insert citation*. The first time you begin working on any document, you will be prompted to select the citation style, as shown in Figure 6.

Once you have selected the citation style, the item that you want to cite is searched as selected from the search box (shown in Figure 7). Based on the citation style selected, either the numbers or authors' names will appear in the document.

You can insert the bibliography by clicking the *Zotero – Insert Bibliography* icon from the tool bar. The citation style is modified by clicking the *Document Preferences* icon. The citations can be deleted or modified as well. All the subsequent changes are automatically managed by Zotero.

## Zotero Sync

Zotero provides the feature of syncing your library across multiple systems. The collection of items in your library is stored on the Zotero server. You need to log in and click on the *Sync with Zotero server* icon to perform the sync operation, which can be done from either the standalone or the add-on version.

## Zotero groups

Zotero groups allow you to collaborate with people of similar interests and share resources with them. Groups facilitate sharing, collaboration and discovery among the members. The groups are of three types: public, open membership; public, closed membership; and private membership. The group settings can be chosen based on needs. Zotero allows modification of these settings. While some reference managers restrict the group management features in their free version, offering more options only in their premium versions, Zotero doesn't have any such restrictions. This becomes a significant factor while choosing Zotero as your favourite reference manager.

## Zotero for developers

Developers can contribute towards enriching Zotero further in more than one way. The contributions could be:

- To develop a plugin to introduce new functionality. Detailed information about how to develop plugins can be gathered from [https://www.zotero.org/support/dev/sample\\_plugin](https://www.zotero.org/support/dev/sample_plugin).
- To create a new citation style using Citation Style Language (CSL). Further information on this is available at [https://www.zotero.org/support/dev/citation\\_styles](https://www.zotero.org/support/dev/citation_styles)
- To use the Zotero API in different applications to extend the utility of Zotero.
- To improve the Zotero client by bug-fixing or providing more features.
- To write Zotero translators that would improve the import

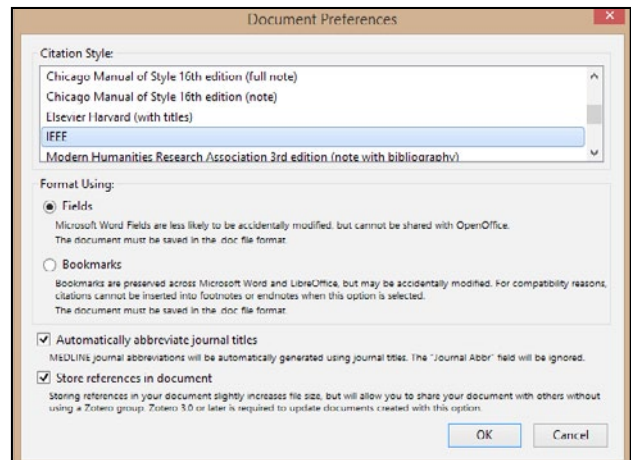


Figure 6 Inserting a citation in Zotero

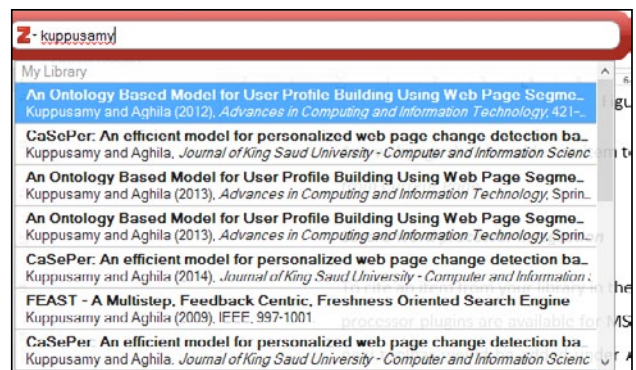


Figure 7: Zotero citation-selection style

and export of metadata items from a variety of sources.

Non-developers can also contribute to Zotero by testing the tool, translating it to other languages, adapting it at their workplace or by simply spreading the word about it, which would be a great help for both the end users and the open source community.



## Acknowledgement

The author would like to acknowledge the support and encouragement offered by Dr G. Aghila, professor and head, Department of Computer Science & Engineering, National Institute of Technology Puducherry, Karaikkal, in promoting the adoption of open source tools like Zotero among research scholars and students of higher classes.

## References

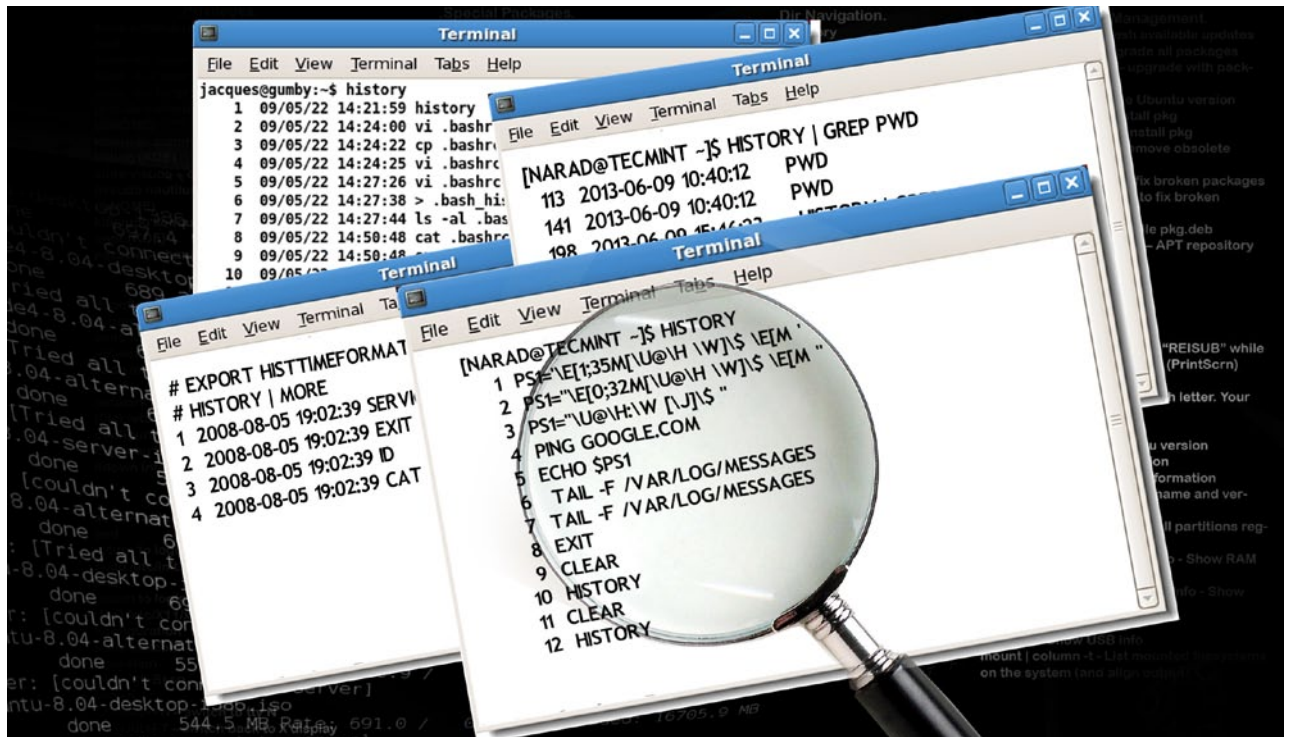
- [1] Detailed Comparison of Reference Manager Software : [https://workspace.imperial.ac.uk/library/Public/Reference\\_management\\_software\\_comparison.pdf](https://workspace.imperial.ac.uk/library/Public/Reference_management_software_comparison.pdf)
- [2] Zotero Forums: <https://forums.zotero.org/categories/>

## By: Dr K S Kuppusamy

The author is assistant professor of computer science, School of Engineering and Technology, Pondicherry Central University. His research interests include Web information retrieval, mobile computing and accessible computing. He can be reached via mail at [kskuppup@gmail.com](mailto:kskuppup@gmail.com)



# Use the History Command Effectively to Make Your Work Easier



Linux users have to make use of the command line at one time or the other. The History command saves them from repeatedly typing certain commands. This article will help readers to leverage the History operations available.

While working with the GNU/Linux environment, one is forced to spend some time on the command line. One is likely to repeat common commands often, and typing variations on those commands even more frequently. This can be good practice in the beginning but later it gets annoying and the repetition of commands slows one down.

Luckily, we can save some time with the use of the *History* command and its variations.

In this section, we will look at how we can explore this command. For this demonstration, I am using Ubuntu 14.04 LTS, but the process is the same for most GNU/Linux distributions.

## Using the *History* command

The GNU/Linux History library is able to keep track of all the previously executed commands and this information can be helpful in writing new ones and executing them. The *History* command can be used to list a log of the commands you have typed.

This log is called the 'history'. And the command used to access it is *history*.

```
[bash]$ history
599 cat .bashrc
600 nautilus
601 vi
602 ls -l
```

In the above output, the former parameter is the command number in the History file and the latter is the actual command.

This will list all the previously executed commands.

Typing '*history n*' will list the last '*n*' commands from the log.

The following command will list the last five commands from the history log.

```
[bash]$ history 5
599 cat .bashrc
600 nautilus
601 vi
602 ls -l
603 history
```

The following are some keys to scroll through the history and jump to either end:

*UP* arrow key: Scrolls backwards in history.

*CTRL* + *p*: Scrolls backwards in history.

*DOWN* arrow key: Scrolls forward in history.

*CTRL* + *n*: Scrolls forward in history.

*ALT*+*Shift*+*.*: Jumps to the end of the history.

*ALT*+*Shift*+*,*: Jumps to the beginning of the history.

## Variables configuration

The behaviour of the history log file is controlled by a few variables and these can be found or can be added to the *.bashrc* file. So let us have a look at the *.bashrc* file. This is present in the home directory. To view its contents, open a terminal and fire up the *cat* command followed by the name (i.e., *cat .bashrc*) or simply navigate to the home directory and press *CTRL*+*h* and open it with the text editor. Let's review some of the variables.

**HISTCONTROL:** Now, it would make sense if you keep some of the variables' values as they are, by default, like the value of the variable *HISTCONTROL=ignoreboth* (which will avoid duplication of commands) and *shopt -s histappend* (which will append newly executed commands to the previous history).

Although *HISTCONTROL=ignoreboth* is good enough for us, some of the other possible values for *HISTCONTROL* along with the syntax are:

```
[bash]$ export HISTCONTROL=ignoredups
```

(To eliminate duplicates, set *HISTCONTROL* to *ignoredups*)

```
[bash]$ export HISTCONTROL=erasedups
```

(To eliminate duplicates across the whole history, set the *HISTCONTROL* to *erasedups*)

```
[bash]$ export HISTCONTROL=ignorespace
```

(When you execute a command, you can instruct *history* to ignore the command by setting *HISTCONTROL* to *ignorespace* and typing a space in front of the command)

**HISTSIZE and HISTFILESIZE:** Some of the other variables in the *.bashrc* file are *HISTSIZE* and *HISTFILESIZE*.

- *HISTSIZE* is the number of lines of the command to store in the history 'list' (i.e., in the memory)
- *HISTFILESIZE* is the number of lines of the command to store in history 'file' (i.e., the actual history file)

These variables are set to an integer value. To change these values, use the following command:

```
[bash]$ export HISTSIZE=2000
```

```
[bash]$ export HISTFILESIZE=2000
```

*HISTSIZE* can be used to disable the use of history.

If you want to disable history all together and don't want bash shell to remember the commands you've typed, set the *HISTSIZE* to 0 as shown below:

```
[bash]$ export HISTSIZE=0
```

Note that now history will not display anything.

**HISTFILE:** The *HISTFILE* variable can be used to store the history log to an alternate file. To do so, execute the following command:

```
[bash]$ export HISTFILE=.bash_alternate_history
```

Now the history will be stored in *.bash\_alternate\_history*.

**HISTTIMEFORMAT='%F %T':** Typically, when you type *history* from the command line, it displays the previously executed commands, but sometimes it may be beneficial to display the time stamp along with the command. To do so, execute the following command:

```
[bash]$ export HISTTIMEFORMAT='%F %T '
```

Now, on executing the history command, the execution date and time along with the command will be displayed:

```
[bash]$ history
```

```
16 2014-11-09 17:51:11 ll
```

```
17 2014-11-09 17:54:32 cd ..
```

```
13 2014-11-09 17:57:37 pwd
```

```
14 2014-11-09 17:58:47 history
```

**HISTIGNORE:** Sometimes you may not want your history to remember some basic commands such as *ls*. Use *HISTIGNORE* to specify all the commands that you want to ignore from the history. Please note that adding *ls* to the *HISTIGNORE* ignores only *ls* and not *ls -l*. So, we have to be specific about the command that we would like to ignore from the history. Thus, executing the following command will ignore *ls* and *ls -l* commands:

```
[bash]$ export HISTIGNORE="ls:ls -l:"
```



**Note: 1)** Sometimes you may be required to re-login for the changes to take effect.

2) To make permanent changes, add the variables with their values in the *.bashrc* file and re-login for the changes to take effect.

## Event designators for history

Consider the history log for the upcoming event designator examples:

```
[bash]$ history
```

```
1 pwd
```

```
2 ls
```

```
3 cd Documents/
```

```
4 ll
```

```
5 nautilus
```

```
6 history
```

```
!n
```

The above executes command number ‘n’.

```
[bash]$ !1
```

```
pwd
```

```
/home/christmas/Documents
```

```
!-n
```

This executes current command minus ‘n’.

```
[bash]$ !-3
```

```
total 1648
```

```
drwxr-xr-x 5 christmas christmas 4096 Nov 6 11:48 ./
```

```
drwxr-xr-x 35 christmas christmas 4096 Nov 8 18:34 ../
```

```
-rw-rw-r-- 1 christmas christmas 8164 Nov 4 19:05 24914.png
```

```
-rw-rw-r-- 1 christmas christmas 27817 Nov 5 12:05  
400_F_37625904_mN1lKUnYH3cnVRKPlankmHcz1zw1SjN.jpg
```

```
!!
```

This executes the previous command. This is similar to ‘!-1’.

```
[bash]$ !!
```

```
1 pwd
```

```
2 ls
```

```
3 cd Documents/
```

```
4 ll
```

```
5 nautilus
```

```
6 history
```

```
!string
```

...executes the most recent command preceding the current position in the history list, starting with the specified string.

```
[bash]$ !pw
```

```
pwd
```

```
/home
```

```
^string1^string2^
```

Quick substitution is of the form `^original^replacement^`. Repeat the last command, replacing the *original* with the *replacement*.

```
[bash]$ cd documents
```

```
bash: cd: documents: No such file or directory
```

```
[bash]$ ^documents^Documents^
```

```
cd Documents
```

```
[bash]$
```

```
Ctrl+r
```

Pressing *Ctrl+r* from the command prompt will display the *reverse-i-search* prompt as shown below.

Search previously executed commands with the *Ctrl+r* command. Once you’ve found the command you’re looking for, press *Enter* to execute it.

```
[bash]$ cat test
```

```
This is a test file.
```

```
(reverse-i-search)`test': cat test [Enter]
```

```
[bash]$ cat test
```

```
This is a test file.
```



## Browsing through history

To find all commands that involve a certain string, simply pipe it to *grep*.

```
history | grep cd
```

```
[bash]$ history | grep cd
```

```
3 cd Documents/
```

If you wish to view the history one page at a time, you can use the command below. Now, you can simply use the space bar to view one page at a time or use the down arrow to view one line at a time:

```
history | less
```

To view just the last 10 commands, you can use the following:

```
history | tail
```

*tail* when used with a *-n* option displays only the last *n* lines.

```
[bash]$ history | tail -5
```

```
29 man tail
```

```
30 man grep
```

```
31 nautilus
```

```
32 history | less
```

```
33 history | tail -5
```

## History command line options

```
[bash]$ history
```

```
1 pwd
```

```
2 ls
```

```
3 cd Documents/
```

```
4 ll
```

```
5 nautilus
```

```
6 history
```

*-c*

The above code clears the history log.  
syntax: *history -c*

```
[bash]$ history
```

```
[bash]$
```

*-d -d offset*

The code above deletes the history entry at position *offset*.  
syntax: *history -d 5* (*Deletes the entry at position 5*)

```
[bash]$ history
```

```
1 pwd
```

```
2 ls
```

```
3 cd Documents/
```

```
4 ll
```

```
5 history
```

*-a*

The above code appends the new history lines (entries in the current session) to the history file.  
syntax: *history -a*

*-r*

The command line given above reads the current history file.  
syntax: *history -r*


*-w*

The above command line writes out the current history to the specified file. This option is useful to export the contents of the history file.

syntax: *history -w alternate\_histfile.txt*

*-s*

The arguments are added to the end of the history list as a single entry, as follows.

syntax: *history -s argument.* 

**By: Chetan Tarale**

The author is interested in the security domain. He likes to explore GNU/Linux and emerging open source technologies.

# Tweak Firefox to Get a Better Browsing Experience



Here are a few simple tricks to improve Mozilla Firefox's performance and privacy settings. You can also enjoy ads-free browsing.

**H**ave you ever tried entering *about:config* in a Firefox tab and checked the available configuration fine-tuning options? If not, try it to customise your Firefox, but be careful while you do these changes. Also, restart the browser for the changes to take effect.

## Improve speed

You can improve the browsing speed in Firefox in quite a few ways, some of which are listed below.

**Enable pipelining:** Start by entering *about:config* in a new Firefox tab. Click on the 'I'll be careful, I promise!' button and search for 'pipelining' (without quotes). Double click on the desired option and ensure the following settings:

```
network.http.version      1.1
network.http.pipelining   true
network.http.pipelining.maxrequests 8
network.http.pipelining.ssl true
network.http.proxy.pipelining true
```

By applying these changes, Firefox is configured to enable the pipelining feature of the http 1.1 protocol and send eight simultaneous http requests to the website. This increases your browsing speed, provided your Internet connection is fast enough to handle these simultaneous requests. Going beyond eight will not help since it is the maximum limit for Firefox.

**Disable animation and placeholder images:** You can do this by changing the following options.

```
browser.tabs.animate      false
browser.panorama.animate_zoom false
browser.display.show_image_placeholders false
```

**Disable 'safe browsing' (at your own risk):** Change the following option.

```
browser.safebrowsing.enabled false
```

**Run Firefox in RAM to speed it up:** Follow the instructions in <http://www.wikihow.com/Speed-Up-Firefox-by-Running-It-In-RAM>

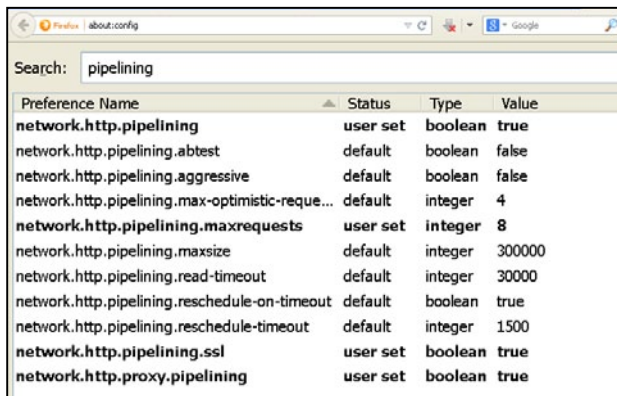


Figure 1: Pipelining options

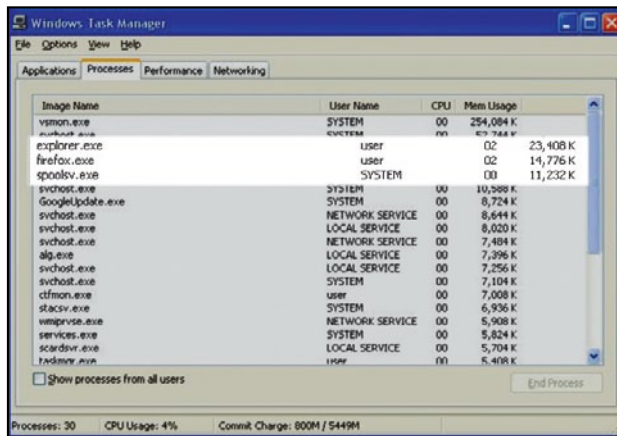


Figure 2: Firefox memory use – when minimised

**Configure Firefox to use 10 MB RAM when minimised:** Open the *about:config* box; right-click anywhere in the 'Preference name' pane, then create a new Boolean variable – *config.trim\_on\_minimize*, and set its value to 'True'.

This is especially useful for systems with low RAM, or if you see heavy RAM usage by Firefox under the task manager.

## View source code in your preferred editor

Start by entering *about:config* in a new Firefox tab, search for *view\_source.editor* and change the following options:

*view\_source.editor.external*      true  
*view\_source.editor.path*      Enter your editor path

Now, pressing *Ctrl-U* will open the source code in the editor mentioned in the path. For example, you may enter *c:\windows\notepad.exe* for viewing (or editing) the code in Notepad. This could really be helpful for developers and testers.

## Ensure privacy while browsing

**Disable third party cookies, and delete all cookies when you close Firefox:**

Tools - Options - Privacy  
 Use custom settings for history  
 Accept cookies from sites  
 Accept third party cookies - never

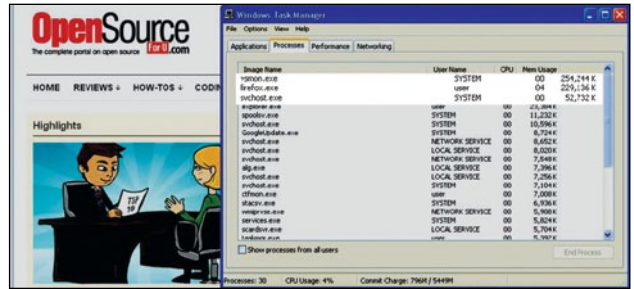


Figure 3: Firefox memory use – when maximised



Figure 4: Trackers blocked by Ghostery on a prominent media website

Keep until - I close Firefox  
 Select 'Tell sites that I do not want to be tracked' under Tools - Options - Privacy

Websites store information about you in cookies, which are accessed whenever you visit the same site again. This is primarily used for various reasons such as 'Keep me signed in', providing customised contents to you.

Third party cookies are those that are placed in your computer by sites that you are not visiting—you may wish to block them all together.

**Use the add-on, Ghostery:** Most of the websites you browse through, track and record your activity – even though you have selected 'Tell sites that I do not want to be tracked'. They could also be selling this information. Ghostery detects these trackers and allows you to block them. It's available for free at <https://addons.mozilla.org/en-US/firefox/addon/ghostery/>

## Use the Adblock Plus add-on to block advertisements

On many websites, advertisements are one of the the biggest nuisances. Adblock Plus blocks banners, pop-ups and video advertisements. Install it to see the difference—it is unbelievable.

There are various user configurable settings for this add-on:

1. Subscribe to various filters to block advertisements.
2. Disable this add-on for particular websites
3. By default, this add-on allows some non-intrusive advertisements such as those by google to be displayed - user can disable them from filter preferences.
4. Create custom rule for blocking particular contents such as advertisement, banner or video if not automatically blocked. **END**

**By: Rajesh Deodhar**

The author is an IS auditor, network security consultant and trainer. He is an industrial electronics engineer with CISA, CISSP and DCL certifications. Please feel free to contact him on [rajesh@omegasystems.co.in](mailto:rajesh@omegasystems.co.in)



# Why Care About **Technical Debt**?



Quick fix solutions in software design, if not corrected, can lead to ‘technical debt’, which if not addressed in time, can lead to ‘technical bankruptcy’. This article, which is an extract from the first chapter of the book, ‘Refactoring for Software Design Smells: Managing Technical Debt’, looks into the causes and impact of technical debt, and gives a few tips on how to manage it.

**T**echnical debt is the debt that accrues when you knowingly or unknowingly make wrong or non-optimal design decisions.

Technical debt is a term coined by Ward Cunningham in a 1992 report. It is analogous to financial debt. When a person takes a loan (or uses his credit card), he incurs debt. If he regularly pays the installments (or the credit card bill) then the created debt is repaid and does not lead to further problems. However, if the person does not pay his installments (or bills), a penalty in the form of interest is applicable and this mounts every time he misses making a payment. In case the person is not able to pay the installments (or bills) for a long time, the accrued interest can make the total debt so ominously large that the person may have to declare bankruptcy.

Along the same lines, when software developers opt for a quick fix rather than a proper well-designed solution, they introduce technical debt. It is okay if the developers pay back the debt on time. However, if they choose not to or forget about the debt created, the accrued interest on the technical debt piles up, just like financial debt. The debt keeps increasing over time with each change to the software; thus, the later the developers pay off the debt, the more expensive it is to pay off. If the debt is not paid at all, then eventually, the pile-up is so huge that it becomes immensely difficult to change the software. In extreme

cases, the accumulated technical debt is too big to ever be paid off and the product has to be abandoned. Such a situation is called technical bankruptcy.

## What constitutes technical debt?

There are multiple sources of technical debt (see Figure 1). Some of its well-known dimensions include (with examples):

- *Code debt:* Static analysis tool violations and inconsistent coding style.
- *Design debt:* Design smells and violations of design rules.
- *Test debt:* Lack of tests, inadequate test coverage, and improper test design.
- *Documentation debt:* No documentation for important concerns, poor documentation and outdated documentation.

This book is primarily concerned with the design aspects of technical debt, i.e., design debt. In other words, when the author refers to technical debt in this book, he implies design debt.

To better understand design debt, let us take the case of a medium-sized organisation that develops software products. To be able to compete with other organisations in the market, the former obviously needs to release newer products into the market faster and at reduced costs. But how does this impact its software development process?

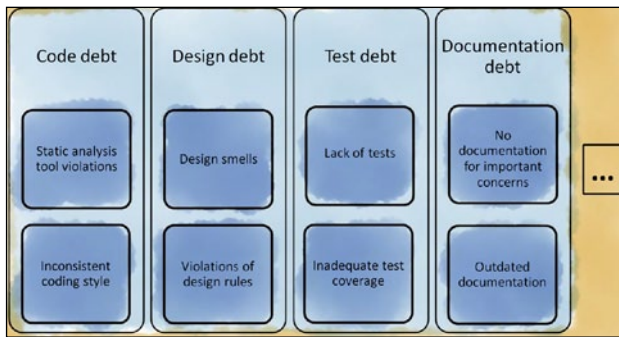


Figure 1: Dimensions of technical debt

As one can imagine, its software developers are expected to implement features faster. In such a case, the developers may not have the opportunity or time to properly assess the impact of their design decisions. As a result, over time, such a collection of individual localised design decisions starts to degrade the structural quality of the software products, thereby contributing to the accumulation of design debt.

If such a product were to be developed just once and then no longer maintained, the structural quality would not matter. However, most products are in the market for a long time period and therefore have an extended development and maintenance life cycle. In such cases, the poor structural quality of the software will significantly increase the effort and time required to understand and maintain the software. This will eventually hurt the organisation's interests. Thus, it is extremely important for organisations to monitor and address the structural quality of the software. The work that needs to be invested in the future, to address the current structural quality issues in the software, is design debt.

An interesting question in the context of what constitutes technical debt is whether defects/bugs are a part of this debt. Some argue that defects (at least some of them) originate due to technical debt, and thus are part of technical debt. Their view is that if managers decide to release a software version despite it having many known yet-to-be-fixed defects, these are a part of the technical debt that has been incurred.

However, there are others in the community who argue that defects do not constitute technical debt. Their view is that the main difference between defects and technical debt is that defects are visible to the users while technical debt is largely invisible. We support this stance. In our experience, defects are rarely ignored by the organisation and receive much attention from the development teams. On the other hand, issues leading to technical debt are mostly invisible and tend to receive little or no attention from the development teams. Why does this happen?

This happens because defects directly impact the external quality attributes of the software that are directly visible to the end users. Technical debt, on the other hand, impacts the internal quality of the software system, and is not directly perceivable by the end users of the software.

Organisations value their end users and cannot afford to lose them; thus, defects get the utmost attention while issues related to 'invisible' technical debt are usually deferred or ignored. Thus, from a practical viewpoint, it is better to leave defects out of the umbrella term 'technical debt', so that they can be dealt with separately; otherwise, one would fix defects and mistakenly think that the technical debt has been addressed.

## What is the impact of technical debt?

Why is it important for a software practitioner to be aware of technical debt and keep it under control? To understand this, let us first understand the components of technical debt, which is a result of the 'principal' (the original hack or shortcut), and the accumulated interest incurred when the principal is not fixed. The interest component is compounding in nature; the more you ignore or postpone it, the bigger the debt becomes over time. Thus, it is the interest component that makes technical debt a significant problem.

Why is the interest compounding in nature for technical debt? One major reason is that often, new changes introduced in the software become interwoven with the debt-ridden design structure, further increasing the debt. Then, going forward, when the original debt remains unpaid, it encourages or even forces developers to use 'hacks' while making changes, which further compounds the debt.

Jim Highsmith describes how the 'Cost of Change' (CoC) varies with technical debt. A well-maintained software system's actual CoC is close to the optimal CoC; however, with the increase in technical debt, the actual CoC also increases. As previously mentioned, in extreme cases, the CoC can become prohibitively high leading to technical bankruptcy.

Apart from technical challenges, technical debt also impacts the morale and motivation of the development team. As technical debt mounts, it becomes difficult to introduce changes and the team involved with development starts to feel frustrated and annoyed. The frustration is further compounded because the alternative—i.e., repaying the whole technical debt—is not a trivial task that can be accomplished overnight.

It is believed that technical debt is the reason behind software faults in a number of applications across domains, including financing. In fact, a BBC report clearly mentions technical debt as the main reason behind the computer controlled trading error that decimated the balance sheet of Knight Capital, an American global financial services firm in 2012.

## What causes technical debt?

The previous section discussed the impact of technical debt in a software system. To pay off the technical debt or to prevent a software system from accruing technical debt, it is

important to first think about why technical debt happens in the first place.

Ultimately, the decisions made by a manager, architect, or developer introduce technical debt in a software system. For instance, when managers create or modify a project plan, they can decide whether to squeeze in more features in a given time span or to allocate time for tasks such as design reviews and refactoring that can ensure high design quality. Similarly, an architect and a developer have to make numerous technical decisions when designing or implementing the system. These design or code-level decisions may introduce technical debt.

Now, the question is: Why do managers, architects or developers make the decisions that introduce technical debt in the software system? In addition to lack of awareness about technical debt, the software engineering community has identified several common causes that lead to technical debt, such as:

- **Schedule pressure:** Often, while working under deadline pressures to 'get-the-work-done' as soon as possible, programmers resort to hasty changes. For example, they embrace 'copy-paste programming', which helps get the work done. They think that as long as there is nothing wrong syntactically and the solution implements the desired functionality, it is an acceptable approach. However, when such code duplication accumulates, the design becomes incomprehensible and brittle. Thus, a tight schedule for the release of a product with new features can result in a product that has all the desired features but has incurred huge technical debt.
- **Lack of good/skilled designers:** Fred Brooks, in his classic book 'The Mythical Man Month', stressed the importance of good designers for a successful project. If designers lack understanding of the fundamentals of software design and principles, their designs will lack quality. They will also do a poor job while reviewing their team's designs and end up mentoring their teams into following the wrong practices.
- **Not adequately applying design principles:** Developers without the awareness or experience of actually applying sound design principles often end up writing code that is difficult to extend or modify.
- **Lack of awareness about design smells and refactoring:** Many developers are unaware of design smells that may creep into the design over time. These are indicative of poor structural quality and contribute to technical debt.


Design smells can be addressed by timely refactoring. However, when developers lack awareness of refactoring and do not perform it, the technical debt accumulates over time. Often, given the different cost and schedule constraints of a project, it may be acceptable to temporarily incur some technical debt. However, it is critical to pay off the debt as early as possible.

## How to manage technical debt

It is impossible to avoid technical debt in a software system; however, it is possible to manage it. This section provides a brief overview of high-level steps required to manage technical debt.

**Increasing awareness of technical debt:** Awareness is the first step toward managing technical debt. This includes awareness of the concept of technical debt, its different forms, the impact of technical debt, and the factors that contribute to it. Awareness of these concepts will help your organisation take well-informed decisions to achieve both project goals and quality goals.

**Detecting and repaying technical debt:** The next step is to determine the extent of technical debt in the software product. Identifying specific instances of debt and their impact helps prepare a systematic plan to recover from the debt. These two practical aspects of managing technical debt are addressed in detail in Chapter 8 of the book.

**Prevent accumulation of technical debt:** Once technical debt is under control, all concerned stakeholders must take steps to ensure that the technical debt does not increase and remains manageable in the future. To achieve this, the stakeholders must collectively track and monitor the debt and periodically repay it to keep it under control. 



**Note:** 'Refactoring for Software Design Smells: Managing Technical Debt', published by Morgan Kaufmann/Elsevier, is available worldwide from November 2014 (URL: <http://amzn.com/0128013974>). This extract is published with permission from Morgan Kaufmann/Elsevier, MA, USA.

## References

- [1] Cunningham W, 'The WyCash portfolio management system'. In: Addendum to the proceedings on object-oriented programming systems, languages, and applications. OOPSLA '92; 1992.
- [2] Highsmith J, 'Zen and the art of software quality'. In: Agile2009 conference; 2009.
- [3] Kelion L, 'Why banks are likely to face more software glitches in 2013'. Available at: <http://www.bbc.co.uk/news/technology-21280943>.
- [4] Brooks Jr F P, 'The Mythical Man-month'. Anniversary ed. Boston, MA, USA: Addison-Wesley Longman Publishing Co Inc; 1995./

## By: Girish Suryanarayana Ganesh Samarthayam and Tushar Sharma


Girish Suryanarayana is a senior research scientist at Research and Technology Centre, Siemens Technologies and Services Pvt Ltd, Bengaluru.

Ganesh Samarthayam is an independent consultant and corporate trainer based in Bengaluru. He has contributed numerous articles to LFY/OSFY since 2006.

Tushar Sharma is a technical expert at Research and Technology Centre, Siemens Technologies and Services Pvt Ltd, Bengaluru.



# Develop a GNU/Linux-like OS for a Single Board Computer



Single board computers (SBCs) have a variety of applications such as monitoring and observation devices, kiosk terminals, etc. In this column, the author shares his experience of developing a monolithic-kernel GNU/Linux-like OS on the ARM platform for a SBC.

**S**ingle board computers (SBCs) have become pretty popular in a wide variety of fields. As the core component of computer systems as well as of embedded systems, the operating system plays a very important role in these systems.

For the purpose of technical research and teaching a curriculum, I have developed a monolithic-kernel SBC GNU/Linux-like OS on the ARM platform. The article covers a boot loader design called U-boot, building the kernel – uImage, the design of the root file system and the Init process. The single board computer OS (SBC OS) is developed on the Linux platform with the GNU tool chain. The system mainly focuses on helping students to learn about and design tiny operating systems on the ARM platform from scratch when the source code is provided.

## Architecture of the SBC OS

At the top of the SBC OS is the user or application space where user applications are executed. Below the user space is the kernel space, where the SBC OS kernel resides.

The SBC OS also contains a GNU C library (*glibc*), which provides the system call interface that connects to the SBC OS kernel and provides the mechanism for the transition between the user or application space and the SBC OS kernel. This is important because the kernel and the user application occupy different protected address

spaces. While each user or application space process occupies its own virtual address space, the SBC OS kernel occupies a single address space.

The SBC OS kernel can be further divided into three levels. At the top is the system call interface, which implements the basic functions such as read and write. Below the system call interface is the SBC OS kernel code, which can be more accurately defined as the architecture-independent kernel code. This code is common to all of the processor architectures supported by the SBC OS. Below this is the architecture-dependent code, which forms what is more commonly called a BSP (board support package). This code serves as the processor and platform-specific code for the given architecture.

## Design and implementation

### The U-boot boot loader design

U-boot is an open source, cross-platform boot loader that provides out-of-the-box support for hundreds of SBCs and many CPUs, including PowerPC, ARM, XScale, MIPS, Coldfire, NIOS, Microblaze and x86.

The SBC OS normally resides in large-capacity devices such as hard disks, CD-ROMs, USB disks, network servers and other permanent storage media. When the processor is powered on, the memory does

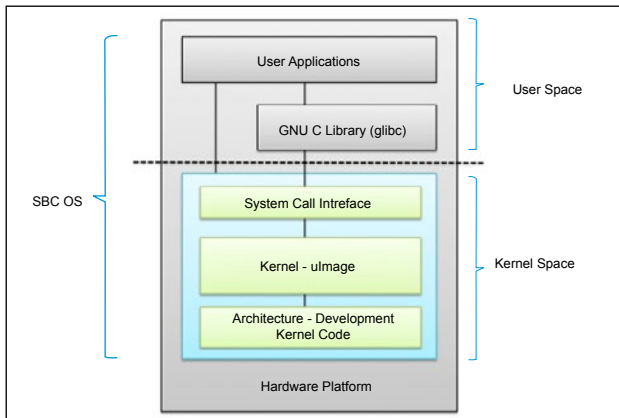


Figure 1: Architecture of SBC OS




Figure 2: Storage of the boot loader, kernel and root file system in the ROM memory (Flash)

not hold an operating system, so special software is needed to bring the SBC OS into the memory from the media on which it resides. This software is normally a small piece of code called the boot loader. On a desktop PC, the boot loader resides on the master boot record (MBR) of the hard drive and is executed after the PC's basic input output system (BIOS) performs system initialisation tasks. In a SBC, the boot loader's role is more complicated because these systems rarely have a BIOS to perform initial system configuration.

At the very minimum, a boot loader for the SBC OS performs these functions:

- Initialises the hardware, especially the memory controller
- Provides boot parameters for the SBC OS
- Starts the SBC OS

## Compiling the boot loader

Boot loader compilation for the SBC OS can be done using the following command:

```
$ make clean && make cubietruck CROSS_COMPILE=arm-linux-gnueabi-
```

## Building the kernel – ulmage

The kernel is the most fundamental software component of the SBC OS. It is responsible for managing the bare hardware within the chosen target system and bringing order to what would otherwise be a chaotic struggle between each of the many different software components on a typical system.

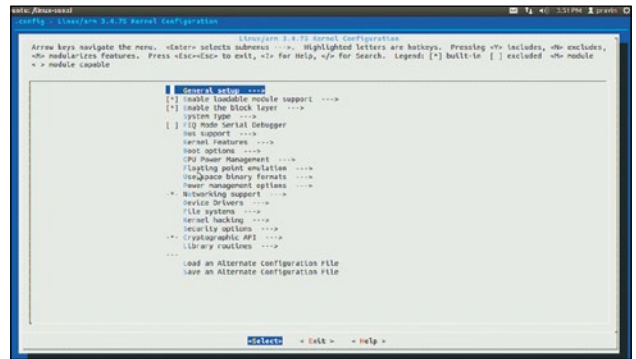


Figure 3: Snapshot of Menuconfig

## Configuring the kernel

The options that we need in the kernel have to be configured before building it. The target is to have an appropriate `.config` file in our Kernel source distribution. Depending on our target, the option menus available will change, as will their content. Some options, however, will be available no matter which embedded architecture we choose. After the environmental set-up, `make menuconfig` runs a text-based menu interface as shown in Figure 3.

```
$ make ARCH=arm CROSS_COMPILE=arm-linux-gnueabi- menuconfig
```

<\*> indicates that the feature is on

<M> indicates that the feature is configured as module

< > indicates that the feature is off

## Compiling the kernel and modules

After saving the kernel configuration in the kernel root directory, our main goals are to compile the ulmage compressed kernel and ulmage kernel modules for the SBC OS using the following commands:

```
$ make ARCH=arm CROSS_COMPILE=arm-linux-gnueabi- uImage -j4
$ make ARCH=arm CROSS_COMPILE=arm-linux-gnueabi- modules -j4
```

## Installing modules

Before we install and boot from our new SBC OS kernel, we should put the new kernel modules in `/lib/modules` with the following command:

```
$ make ARCH=arm CROSS_COMPILE=arm-linux-gnueabi- modules_
install
```

Our new modules should appear in `/lib/modules/version`, where the version is the new kernel version of the SBC OS.

## Designing the root file system

One of the last operations conducted by the Linux kernel during system start-up is mounting the root file system. The Linux kernel itself doesn't dictate any file system structure, but user space applications do expect to find files with

```

ubuntu:/fedora/kernel-source
LD      net/built-in.o
LD      vmlinux.o
MODPOST vmlinux.o
GEN      .version
CHK      include/generated/compile.h
UPD      include/generated/compile.h
CC      init/version.o
LD      init/built-in.o
LD      .tmp_vmlinux1
KSYM     .tmp_kallsyms1.o
AS       .tmp_kallsyms1.o
LD      .tmp_vmlinux2
KSYM     .tmp_kallsyms2.o
AS       .tmp_kallsyms2.o
LD      vmlinux
SYSMAP   System.map
SYSMAP   .tmp_System.map
OBJCOPY arch/arm/boot/Image
Kernel: arch/arm/boot/Image is ready
AS       arch/arm/boot/compressed/head.o
GZIP     arch/arm/boot/compressed/plggy.gz
CC       arch/arm/boot/compressed/misc.o
CC       arch/arm/boot/compressed/decompress.o
CC       arch/arm/boot/compressed/string.o
SHIPPED  arch/arm/boot/compressed/libbifuncs.o
SHIPPED  arch/arm/boot/compressed/ashld3.o
AS       arch/arm/boot/compressed/plggy.gz
AS       arch/arm/boot/compressed/libbifuncs.o
AS       arch/arm/boot/compressed/ashld3.o
LD      arch/arm/boot/compressed/vmlinux
OBJCOPY arch/arm/boot/zImage
Kernel: arch/arm/boot/zImage is ready
IMAGE   arch/arm/boot/uImage
Image Name: Linux-3.4.75
Created: Mon Mar 17 20:22:29 2014
Image Type: ARM Linux Kernel Image (uncompressed)
Data Size: 4616400 Bytes = 4500.20 kB = 4.40 MB
Load Address: 40000000
Entry Point: 40000000
Image arch/arm/boot/uImage is ready
root@ubuntu:/fedora/kernel-source#

```

Figure 4: Snapshot of SBC OS kernel compilation - ulmage

specific names in specific directory structures. Therefore, it is useful to follow the de facto standards that have emerged in Linux systems.

```
$ debotstrap --no-check-gpg --arch=armhf --foreign wheezy
```

## The Init process and runlevels

In conventional Linux systems, Init is the first process started when a Linux kernel boots and it's the ancestor of all processes. Its primary role is to start appropriate service processes for the 'state' the system is to run in at boot and to shut down or start appropriate services if the system state changes (such as changing to the halt/shut down state). It can also create consoles and respond to certain types of events.

Init's behaviour is determined by its configuration file /etc/inittab. Lines in /etc/inittab have the following syntax:

```
id:runlevels:action:process
```

where:

```

id      - 1-4 (usually 2) character name for the line,
totally arbitrary;
runlevels - a list of runlevels the line applies to;
action   - what init is to do and/or under what
conditions;
process  - program/command to be run.

```

Here are the typical runlevels and what they mean for Red Hat family distros:

```

0 - halt system
1 - single user mode (no GUI)
2 - multiuser mode, no networking (no GUI)
3 - multiuser mode, networking (no GUI)
4 - unused

```




Figure 5: The Cubietruck SBC

```

5 - multiuser mode (GUI/X11) //FUTURE WORK
6 - reboot system

```

## A demo of the SBC OS and the single board computer called Cubietruck

FOSSOF 1.0 (Free and Open Source SOFTware) is a SBC OS that can be obtained from <https://github.com/gselvapravin/FossoF>. Developed by yours truly, it can be cloned by using the following commands:


```

$ sudo apt-get -y install git
$ cd ~
$ git clone https://github.com/gselvapravin/FossoF
$ chmod +x ./FossoF/fossof.sh
$ cd ./FossoF
$ ./fossof.sh

```

The compiled image will be located in /tmp/FossoF/output/debian\_rootfs.raw.gz. To write it to an SD card, decompress it and use Image Writer (Windows) or DD-it in Linux by using the following command:

```
$ dd bs=1M if FossoF_x.x_vga.raw of=/dev/sdx
```

Cubietruck is an SBC and is the third board of Cubieteam; so it is also called Cubieboard3. It's a new PCB model adopted with the Allwinner A20 main chip, just like Cubieboard2. But it is enhanced with some features, such as 2GB memory, an on-board VGA display interface, 1000M network interface, Wi-Fi+BT on board, support for Li batteries and RTC, and the SPDIF audio interface. **END** 

### By: Pravin Selva

The author is an Open Source enthusiast who wishes to share his knowledge in the true spirit of Open Source, and a Programmer Analyst Trainee at Cognizant Technology Solutions India Private Limited ("Cognizant").



# Install BURG, the Brand-new Universal loader from GRUB

There's a new boot loader on the block called BURG. Install it and discover how to enjoy it with this simple tutorial!

If you run multiple OSs on your machine, GRUB is probably the first screen you will look at when you turn ON your system. GRUB stands for GRand Unified Boot loader, and is a part of the GNU project. It is the default boot loader that comes with all UNIX-like OSs. GRUB provides a pretty basic menu to choose from—a list of the installed OSs and the entries that come with them, such as recovery modes and memtests.

## So why switch to BURG?

Quite simply, because GRUB is now plain boring. Over the years, the user interface in our desktop environments has improved a lot. It has become more user-friendly and highly customisable. But sadly, GRUB has remained pretty much the same. It still offers the bash-like, command line interface of old!

BURG is a cool replacement to GRUB. It can turn your boot loader into what's shown in Figure 2. It's minimalistic and appealing.

## How to set up BURG

First, you need to add a new repository; so enter the following command in a terminal:

```
sudo add-apt-repository ppa:n-muench/burg -y && sudo apt-get update
```




Figure 1: The GRUB menu




Figure 2: The BURG menu

Then you need to install BURG and some themes, for which you can use the following command:

```
sudo apt-get install burg burg-themes
```

As the installation proceeds, dialogue boxes will pop up for configuration. Please follow the steps carefully.

Once you have successfully installed BURG and its themes, you could disable the memtest option from the boot screen if you want to, by typing the following command:

```
sudo chmod -x /etc/grub.d/20_memtest86+
```




Figure 3: Package configuration

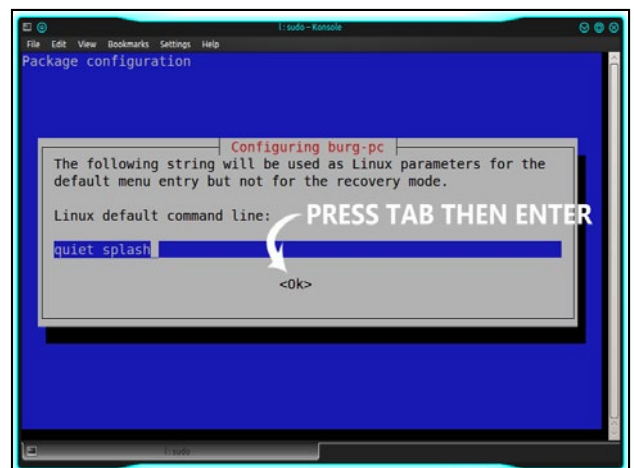


Figure 4: Menu entry options

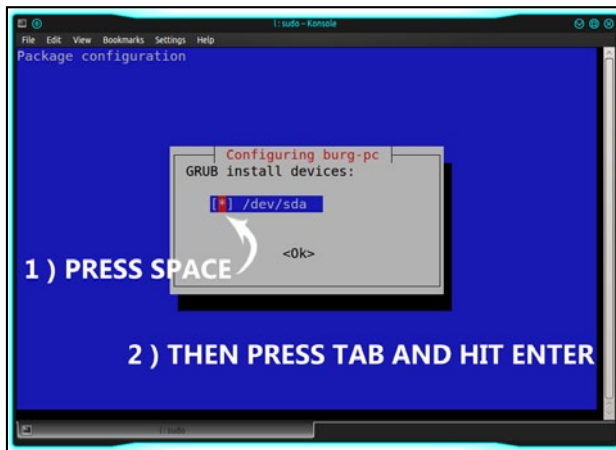


Figure 5: Selecting the device

Now let's see what it looks like. We can emulate the boot screen using the following command:

```
sudo burg-emu
Press
F1 Help      F2 Change Theme    F3 Change Resolution
Arrow-keys to move.
```

For now, just choose your theme, and do not change the resolution. This is because, if your monitor doesn't support the particular resolution you've chosen, the next time you boot your system, you will get a blank screen, which will obviously freak you out. So press F3 during your next boot and choose the resolution by pressing *Enter*. If you get a blank screen while doing so, move up or down to choose another resolution.

I have three OSs installed on my machine (Ubuntu, Windows 8 and Linux Mint). But there are two options each for Ubuntu and Linux Mint. Those are the recovery modes that I mentioned at the beginning, which can be disabled easily. For that, we need to edit the BURG configuration file.




Figure 6: OS options

In this tutorial we will be using a basic editor called Nano. Advanced users may use editors of their choice. Open the terminal and type in the following command:

```
sudo nano /etc/default/burg
```

In this file, we need to edit the following line:

```
#GRUB_DISABLE_LINUX_RECOVERY="true"
```

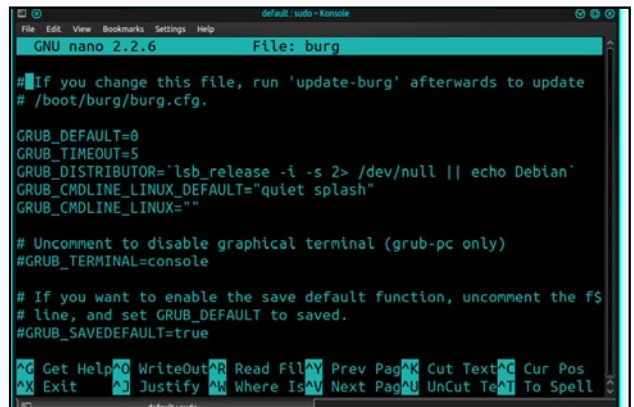


Figure 7: Editing BURG configuration

You need to uncomment the line, by removing the '#' character.

This file also contains the time-out period for the boot screen, if you want to change it. Locate this line and change the R-value, which is in seconds.


```
GRUB_TIMEOUT=5
```

To save the changes made, press *Ctrl-x*, then 'y', and then press 'Enter'.

We now need to configure these changes into BURG and view the modified boot screen. For that, use the following command:

```
sudo update-burg
sudo burg-emu
```

Now the two extra recovery entries would be gone. That's it. Enjoy booting with BURG!

To get more themes, visit <http://www.deviantart.com/> 

## References

- [1] <https://wiki.archlinux.org>
- [2] <http://www.dedoimedo.com>

## By: Chandramohan Sudar

The author loves programming and tweaking Android, Linux and Windows. He blogs about his experiences at <http://chandruscm.wordpress.com/>

# “IT requirements, business constraints and workload requirements determine the best choice for a server”

The latest trends in the mobile, cloud, social networking and Big Data arenas pose new challenges to IT heads. They are forced to develop new approaches towards the data centre and its processing engine - the server, to enable successful business results. To address these challenges, HP recently unveiled its new ProLiant Gen9 servers that aim to help customers reduce costs and complexity, accelerate IT service delivery and enable business growth. IT organisations are under increasing pressure to quickly and efficiently deliver more applications and services, as well as store and deliver exponentially increasing amounts of data while reducing costs, and hence are constrained by traditional data centre and server architectures. HP claims that the new ProLiant Gen9 servers are designed to address these challenges, delivering flexible, scalable computing resources that are aligned to companies' business goals. But this doesn't end the confusion for IT heads. They still need to figure out what the best choice for their workloads and businesses is.

**Diksha P Gupta** from **Open Source For You** spoke to **Gopalswamy Srinivasan**, **category manager, servers, HP India**, about the factors that can help IT heads select the best server for their business. Read on...

**Q** Typically, what is the first step that IT heads take, when it comes to making a decision on choosing the right server, whether it is for a mid-scale firm or a large enterprise?

Basically, it all starts with the workloads that they have to deal with and the segment that they are in—that determines what their needs are. For example, if you look at the mid-market customers—they want a single platform that supports a lot of workloads. And they want it to be reliable, but with the lowest cost of ownership. In such cases, the rack servers are the best



**Gopalswamy Srinivasan**, category manager, servers, HP India

option. If you are specifically thinking of virtualisation, and you want to move to the cloud, then you want to be very flexible and agile in order to deliver. In such situations, one wants more of a converged infrastructure for which a different kind of server like the Blade server is of use. In case you are looking to handle mission critical applications, in which you need almost zero downtime, that's when you should think about an integrity solutions system or a scale-up solution.

Before IT heads really make a choice, they have to evaluate their needs. The ultimate choice also depends on the business elements of cost, risks and the skill sets available internally and what the workloads are. When you think of the costs involved, there are multiple elements to look into including power, cooling, space and what you are trying to optimise for. And finally, it is the SLA and how fast you want to deliver the service.

So, it's a combination of IT requirements, business constraints and then the workload requirements. You have to put together all three and make your choice.



**“Data integrity and availability is critical in today’s business environment. The needs of today’s business environment are addressed by providing a family of smart array controllers and smart storage batteries that allow data migration and provide exceptional data protection, while minimising the long-term costs of that protection.”**

### **Q Blade, tower or rack—which is the best choice for an IT manager?**

For businesses that are new to servers, we recommend entry level rack and tower servers that are simple, easy to deploy and affordable, having been designed for SMB businesses that need the right size for first time workload deployment.

For growing businesses or new growth customers, we recommend the mid-range rack and tower servers, which are optimised with the right balance of storage, performance, efficiency and manageability to address multiple workloads for growing SMB and enterprise businesses.

To customers that have traditionally used IT and are seeking the industry servers of choice, we recommend the enterprise class rack and tower servers, as well as the blade servers, which are designed with flexible choices for compute intensive workloads requiring high system performance, manageability, expansion and security for SMB, enterprise and HPC businesses.

The Compute Strategy that HP has adopted is to focus on providing the right compute for the right workload at the right economics, every time. Different server form factors are positioned as per this strategy. HP’s server portfolio addresses the essential needs of customers who are either new to servers or who are positioned to grow and expand their business.

For organisations requiring the most demanding scale-up workloads, the mission critical rack servers are the best, as they deliver unparalleled scalability, reliability and availability to unleash the power of business data.

### **Q What is the right mix of processor and memory that one needs to look at while choosing the server? How important is the right mix?**

The cloud, mobility, Big Data and security are some of the key mega trends that are generating new business opportunities. This is driving line of business (LOB) executives to deliver new products and services faster, increase operational efficiencies and grow revenue, margins and market share. In order to offer a compute platform to address the different application needs, a right mix of compute capacity and memory is mandatory. The number of CPU cores are growing to offer more compute power, and the memory is becoming bigger and faster to feed those CPU cores so that data seek time and latency can be reduced.

### **Q What kind of server support drives should one look for to achieve the ultimate hard disk performance?**

As data, storage and accessibility requirements grow exponentially, storage solutions need to meet a variety of needs. Data integrity and availability is critical in today’s business environment. The needs of today’s business environment are addressed by providing a family of smart array controllers and smart storage batteries that allow data migration and provide exceptional data protection, while minimising the long-term costs of that protection.

### **Q It is said that network attached storage can provide an alternative to servers in some ways. Do you agree?**


A network attached storage (NAS) appliance offers file-based access using the NFS (Network File System) or CIFS (Common Internet File System) protocols. A server connects to a NFS server system to access a file-based storage for the application that runs on it. Such files can also be accessed by another server connected in the network at a different time. NFS/CIFS systems allow a user on a server to access files over a network much like local storage is accessed.

On the other hand, a server is used to deploy applications like Microsoft Office after installing an operating system, like Microsoft Windows 2008. In NAS, customers can’t deploy any applications. Such servers may have local storage or storage over a network.

### **Q How does one choose the right ports for a server?**

The application’s footprint determines the kind and number of ports required for any server. With the increasing proliferation of virtualisation and the cloud, a lot of compute capacity is needed to process large amounts of data. The tech industry is engaged in extracting contextual information out of Big Data for pragmatic decision making. This calls for consolidating a large number of virtual machines in a physical server and providing a 40 GB network bandwidth per server node to support the high-speed VMs’ east-west traffic, multiple low latency concurrent Vmotion sessions, to consolidate structured and unstructured data, and run the contextual computing applications. Apart from data traffic requirements, a separate dedicated management access is required for the servers to be managed individually or as a group.

### **Q What are the parameters to be kept in mind while choosing an operating system for a server?**

OS selection depends on multiple factors, the most important criterion being the business need. Procurement, deployment, maintenance and upgrade cycles determine the selection of the OS, which impacts cost, employee productivity, ease of use and the learning curve. The business application platform dictates the selection of the operating system. **END** 

# Mozilla Smartphones are Here to Replace Feature Phones!

After a long wait, Mozilla has finally launched Firefox-based smartphones in India. Coming from companies like Spice and Intex, Mozilla's devices are affordable and appeal to the entry-level smartphone users. The company plans to compete in this space only, and aims to offer an experience with a difference. **Diksha P Gupta** from **Open Source For You** spoke to **Dr James Ho, senior director of mobile devices, Mozilla**, about the company's strategy in India and how it plans to work with developers to make the Mozilla OS an even more robust system.



Dr James Ho, senior director of mobile devices, Mozilla

**Q** Don't you think you are late by over a year in bringing the Mozilla phone to India? Android is doing pretty well and with the launch of Android One, we have another ecosystem to look forward to. Android is well accepted in all price ranges and today, one can get an Android phone that is not clunky yet has a great user experience at a price that is equivalent to Mozilla's offering. So, how do you plan to position the Mozilla phone in this market?

I think to make a product successful, it is very important to meet the consumers' expectations. I think we took our time to understand what the users' aspirations are and what they are really looking for from their smartphones. We were able to take those aspirations into account. Frankly speaking, I

think it is never too late because it is a huge industry and is evolving so fast that whenever you come with the right product, you can connect with the consumers.

**Q** Firefox OS was first showcased two years ago at CES and you claimed to be bringing out a US\$ 25 phone at that point in time. Over the past two years, you have a footprint in over 17 countries. Are you satisfied with your performance so far?

We have seen some fabulous responses in the 17 countries that we are present in. We have seen people accepting the Mozilla ecosystem and being a part of its growth. We hope to see the same response in India as well. We are looking forward to expanding our horizons in more countries in the near future. We are also looking forward to launching Mozilla smartphones with more partners in the coming months.

**Q** In the price category that Mozilla's phones are available in, who is your prospective customer? Is it the people who are switching from feature phones to smartphones?

Yes, from our perspective, most of the people picking up the Mozilla phones will be those switching from feature phones to smartphones. At the affordable price point where Mozilla devices are available, most people would want to upgrade from feature phones to a smartphone experience. The cheapest Android phone available in India right now is for Rs 3,000, which is a good 35 per cent higher than Mozilla's phone from Spice. That particular price segment is extremely price sensitive, and comprises the largest chunk of people who will pick up this phone. Having said that, I think smartphone users who are upset that their OSs are not upgraded would also want to try their hands on a Mozilla device. We believe that the Firefox OS, which is coming at this price point, will probably offer a little better experience for such users, who get a stale version of Android at this price point. This device will offer a better seamless experience to those who are already owners of a smartphone.

“At this moment in India, we are targeting only the low-cost smartphone market, which can be considered as the feature phone replacement market. So, any OEM ready to work in that segment is welcome to join the journey.”

**Q** What do you think are going to be the challenges that you will face when it comes to spreading the word about the Firefox OS in India?

Right now, we need to spread awareness about the OS amongst the users. Currently, awareness about the Firefox OS is limited to the Firefox browser users. So we will resort to various means of reaching out to the potential users of the Mozilla Firefox ecosystem in the country. We would not want to restrict ourselves to just Tier 1 cities but would also like to penetrate deeper into the smaller cities and towns with our affordable and user-friendly devices.

**Q** What is Mozilla doing to nurture the developers' ecosystem in India with respect to the Firefox OS?


We already have a community of thousands of developers in

India and they would love to bring in local applications for the OS, which will be more suitable for the Indian market. We would also work with our partners to bring in local Indian content. So in the phone, a user can easily search for local content. With every app that we have on the app store, we try to address the local audience to make the ecosystem friendlier.

**Q** But how do you plan to compete with the ecosystem that Android has already created with respect to the developers and user adoption as well?

We try to compete only on the level of the top 100 chosen applications that the Indian people use. So we will try to partner with the makers of those applications to bring them inside the Mozilla ecosystem.

**Q** Which are the other OEMs you are working with, or plan to work with, in India?

We have lined up a lot of Indian companies who are excited to work on the Mozilla ecosystem. At this moment in India, we are targeting only the low-cost smartphone market, which can be considered as the feature phone replacement market. So, any OEM ready to work in that segment is welcome to join the journey. The Firefox OS is very flexible for OEMs to adopt. So we are looking at more cooperation from OEM partners. **END** 

# India's #1 Website For **Electronics** Engineers Working Across The Globe

Established: **1998** | Registered Users: **300,000+**

## Monthly Statistics

Page Impressions: 750,000+

Unique Techies: 150,000

# electronicsforu.com

If it's electronics, it's here



For any queries, please contact our team at [efyenq@efyindia.com](mailto:efyenq@efyindia.com) OR +91-11-26810601




# It's Good To Have A Choice!

**ASSURED  
GIFT WITH EACH  
SUBSCRIPTION**

**GET  
DISCOUNTS**



**OR**



**Robotics Self  
learning Kit**  
(worth Rs 2500)



**Coffee Maker**  
(Worth Rs 1400)



**Electronics Industry  
Directory 2014-15**  
(Worth Rs 750)



**Arduino Uno R3**  
(worth Rs 1575)



**Toaster Oven**  
(Worth Rs 2000)



**Popcorn Maker**  
(Worth Rs 1100)



**EFY Round  
Neck T-Shirt**  
(worth Rs 300)



**Light  
Following  
Robo**  
(Worth Rs 599)



**OSFY Collared T-Shirt**  
(Worth Rs 500)



**8051 Pocket Programmer**  
(worth Rs 730)

**Terms:-** # These rates are applicable for new subscribers as well as renewal by existing subscribers. # Gift items' photos are representative and actual product may vary # The rates are valid for subscribers within India only.  
# Please allow 4-6 weeks for processing of your subscription. # Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only. \* Replacement will be made if intimation of damaged / non-receipt of copies is received within 30 days of its publication \*\* After three months, if you are not satisfied with the magazine, your balance amount will be returned

# SUBSCRIBE TO YOUR FAVOURITE MAGAZINES NOW!



## Option 1: Assured Gifts with every subscription

| Magazine            | Duration | Issues | Cover Price (₹) | You Pay (₹) | Assured Gift:   |
|---------------------|----------|--------|-----------------|-------------|---|
| Electronics For You | 1 Year   | 12     | 600             | 600         | <input type="checkbox"/> Light Following Robo (Worth Rs 599)        |
|                     | 2 Years  | 24     | 1200            | 1200        | <input type="checkbox"/> 8051 Pocket Programmer (worth Rs 730)      |
|                     | 3 Years  | 36     | 1800            | 1800        | <input type="checkbox"/> Arduino Uno R3 (worth Rs 1575)             |
|                     | 5 Years  | 60     | 3000            | 3000        | <input type="checkbox"/> Robotics Self learning Kit (worth Rs 2500) |
| Electronics Bazaar  | 1 Year   | 12     | 1200            | 1200        | <input type="checkbox"/> EFY Round Neck T-Shirt (worth Rs 300)      |
|                     | 2 Years  | 24     | 2400            | 2400        | <input type="checkbox"/> EYP 2014-15 (Worth Rs 750)                 |
|                     | 3 Years  | 36     | 3600            | 3600        | <input type="checkbox"/> Popcorn Maker (Worth Rs 1100)              |
|                     | 5 Years  | 60     | 6000            | 6000        | <input type="checkbox"/> Toaster Oven (Worth Rs 2000)               |
| Open Source For You | 1 Year   | 12     | 1200            | 1200        | <input type="checkbox"/> OSFY Collared T-Shirt (Worth Rs 500)       |
|                     | 2 Years  | 24     | 2400            | 2400        | <input type="checkbox"/> Popcorn Maker (Worth Rs 1100)              |
|                     | 3 Years  | 36     | 3600            | 3600        | <input type="checkbox"/> Coffee Maker (Worth Rs 1400)               |
|                     | 5 Years  | 60     | 6000            | 6000        | <input type="checkbox"/> Toaster Oven (Worth Rs 2000)               |

## Option 2: Upto 50% Discount Offer

| Magazine            | Duration | Issues | Cover Price (₹) | Discount | You Pay (₹) | You Save |
|---------------------|----------|--------|-----------------|----------|-------------|----------|
| Electronics For You | 1 Year   | 12     | 600             | 20%      | 480         | 120      |
|                     | 2 Years  | 24     | 1200            | 30%      | 840         | 360      |
|                     | 3 Years  | 36     | 1800            | 40%      | 1080        | 720      |
|                     | 5 Years  | 60     | 3000            | 50%      | 1500        | 1500     |
| Electronics Bazaar  | 1 Year   | 12     | 1200            | 20%      | 960         | 240      |
|                     | 2 Years  | 24     | 2400            | 30%      | 1675        | 725      |
|                     | 3 Years  | 36     | 3600            | 40%      | 2155        | 1445     |
|                     | 5 Years  | 60     | 6000            | 50%      | 2995        | 3005     |
| Open Source For You | 1 Year   | 12     | 1200            | 20%      | 960         | 240      |
|                     | 2 Years  | 24     | 2400            | 30%      | 1680        | 720      |
|                     | 3 Years  | 36     | 3600            | 40%      | 2160        | 1440     |
|                     | 5 Years  | 60     | 6000            | 50%      | 3000        | 3000     |

## Introducing EFY's Premium Edition

### Electronics For You PLUS

(With software-packed DVD)

- ☐ 1 Year (12 Issues) ~~1200~~ ₹ 960
- ☐ 2 Yrs (24 Issues) ~~2400~~ ₹ 1800



## Free e-zine Access With Every Subscription

To Subscribe Online,  
Visit: <http://subscribe.efyindia.com>

For Online Renewal,  
Visit: <http://renew.efyindia.com>

- Don't Miss A Single Issue!
- Ensure Regular Supply
- FREE Replacement Policy\*
- Order Now For FREE Home Delivery

\* Replacement will be made if intimation of damaged / non-receipt of copies is received within 30 days of its publication  
\*\* If you are not satisfied with the magazine & services your balance amount will be returned

Name \_\_\_\_\_ Designation \_\_\_\_\_ Organisation \_\_\_\_\_

Mailing Address \_\_\_\_\_ City \_\_\_\_\_

Pin Code \_\_\_\_\_ State \_\_\_\_\_ Phone \_\_\_\_\_ Email \_\_\_\_\_

Subscription No. (for existing subscribers only) \_\_\_\_\_. I would like to subscribe to the above (✓)marked magazine(s) starting with the next

issue. Please find enclosed a sum of Rs \_\_\_\_\_ by DD/MO/crossed cheque\* bearing the No. \_\_\_\_\_ dt. \_\_\_\_\_ in favour of

EFY Enterprises Pvt Ltd, payable at Delhi. (\*Please add Rs 50 on non-metro cheque)


Send this filled-in form  
or its photocopy to :

EFY Enterprises Pvt Ltd D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020  
Ph: 011-26810601-03; Fax: 011-26817563; e-mail: [info@efy.in](mailto:info@efy.in)

**EFYGROUP**  
Technology Drives Us

[www.efy.in](http://www.efy.in)





# TIPS & TRICKS



## Debug your application with the GDB:GNU debugger

The purpose of a debugger such as GDB is to allow us to see what is going on 'inside' a program while it executes, or what the program was doing at the moment it crashed.

We can use GDB to debug programs written in C or C++. GDB is invoked with the shell command *gdb*.

Run the following commands to debug a program:

1. Compile the file using the command:

```
g++ -g xxx.cpp (where xxx is any file)
```

2. Execute the program as follows:

```
gdb ./a.out
```

3. The program will enter into:

```
debugging mode
(gdb)
```

4. Type 'b' (set breakpoint) followed by the function name and press 'Enter'.

```
b main (main function) or b xxx (where xxx is any function name)
```

5. Type 'r' (start the debug program) and press 'Enter', which will display the following:

```
Starting program: /root/a.out
Breakpoint 1, main () at xxx.cpp:7 (7 is line number in the file)
```

6. Type 'n' (step program, proceeding through subroutine calls), till the program exits normally.
7. Other commands that can be used for debugging are:
  - s: step program until it reaches a different source line
  - bt: backtrace, which displays the program stack
  - q: quit debugger

– Sowmya Mitra Attaluri,  
mitra2270@gmail.com



## A command that checks the memory used by each process on the server

We usually check memory utilisation using the *top* command, but it gives the result in a percentage. The following command will give the exact value of the memory used by a process and also sort it according to usage:

```
[root@centos ~]# ps -e -orsz,vsz,args= | sort -n | pr -TW$COLUMNS
```

The following command will give the top 10 commands, according to the non-swapped physical memory that a command has used:

```
[root@centos ~]# ps -e -orsz,vsz,args= | sort -n | pr -TW$COLUMNS | tail -10

10960 54804 /usr/libexec/gdm-user-switch-applet -oaf-activate-iid=OAFIID:GNOME_FastUserSwitchApplet_Factory -oaf-ior-fd=28
11164 43852 /usr/libexec/wnck-applet -oaf-activate-iid=OAFIID:GNOME_Wncklet_Factory -oaf-ior-fd=18
11276 57460 nm-applet -sm-disable
12308 52060 gnome-terminal
12724 45156 gnome-panel
13140 46920 /usr/libexec/clock-applet -oaf-activate-iid=OAFIID:GNOME_ClockApplet_Factory -oaf-ior-fd=34
15208 33640 python /usr/share/system-config-printer/applet.py
17404 67636 /usr/bin/gnote -panel-applet -oaf-activate-iid=OAFIID:GnoteApplet_Factory -oaf-ior-fd=21
18320 29620 /usr/bin/Xorg :0 -nr -verbose -auth /var/run/gdm/auth-for-gdm-oQa9lk/database -nolisten tcp vt1
19304 73780 nautilus
[root@centos ~]#
```

– Shailesh Vaidya,  
skimeer@gmail.com



## Discover the power of the 'history' command

### 1. How to display the time stamp using HISTTIMEFORMAT

Typically, when you type 'history' from command line, it displays the command# and the command. For auditing



purposes, it may be beneficial to display the time stamp along with the command, as shown below. To display the time stamp, you need to use the following code:

```
# export HISTTIMEFORMAT='%F %T '
# history | more

1 2014-08-05 19:22:39 service httpd restart
2 2014-08-05 19:22:39 exit
3 2014-08-05 19:22:39 id
4 2014-08-05 19:22:39 ls
```

## 2. How to execute a previous command that starts with a specific word

Type '!' followed by the starting few letters of the command that you would like to re-execute. In the following example, typing '!ps' and then pressing 'Enter' executed the previous command starting with 'ps', which is '*ps aux | grep yp*'.

```
# !ps

ps aux | grep yp
root    16947  0.0  0.1 36516 1264 ?        Sl   13:10
0:00 ypbind
root    17503  0.0  0.0  4124   740 pts/0    S+   19:19
0:00 grep yp
```

## 3. How to erase duplicates across the whole history using HISTCONTROL

The *ignoredups* shown below removes duplicates only if they are consecutive commands. To eliminate duplicates across the whole history, set the *HISTCONTROL* to *erasedups* as shown below.

```
# export HISTCONTROL=erasedups
# pwd
# service httpd stop
# history | tail -3
38  pwd
39  service httpd stop
40  history | tail -3

# ls -ltr
# service httpd stop
# history | tail -6
35  export HISTCONTROL=erasedups
36  pwd
37  history | tail -3
38  ls -ltr
39  service httpd stop
```

Note that the previous *service httpd stop* after *pwd* got erased.

—Amit Singh,  
[amita123b@gmail.com](mailto:amita123b@gmail.com)



## Command line shortcuts

Here are some tips to speed up our work.

```
cmd1;cmd2
```

The above command will run *cmd1* and then execute *cmd2*.

```
cmd1 && cmd2
```

This will execute *cmd2* if *cmd1* is successful.

```
cmd1 || cmd2
```

The above sequence will run *cmd2* if *cmd1* is not successful.

```
Ctrl+a
```

This will move the control to the beginning of the line in CLI.

```
Ctrl+e
```

The command above will move the control to the end of line in CLI.

```
!ser:p
```

The above command will print the last command starting with 'ser'.

—Sumit Chauhan,  
[sumit1203@gmail.com](mailto:sumit1203@gmail.com)



## Viewing random man pages while logging in through SSH

If you wish to view the random man pages of the Linux command while you log in through SSH, add the following line in the *.bashrc* file:

```
/usr/bin/man $(ls /bin | shuf | head -1)
```

—Ranjithkumar T,  
[ranjith.stc@gmail.com](mailto:ranjith.stc@gmail.com)



## Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in *OSFY* so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at [www.opensourceforu.com](http://www.opensourceforu.com). The sender of each published tip will get a T-shirt.

# DVD OF THE MONTH

Try your hands at some of the great Linux distros available.



**Ubuntu 14.10:** Fast, free and incredibly easy to use, the Ubuntu operating system powers millions of desktop PCs, laptops and servers around the world. It is a lightweight, secure and intuitive OS that is free. It comes with a full office suite that also supports Windows file formats. Try it out to discover what's new in this latest release.

**Ubuntu MATE 14.04:** This is a community-developed Ubuntu-based operating system that beautifully integrates the MATE desktop. This stable, easy-to-use operating system with a configurable desktop environment is ideal for those who want the most out of their desktops, laptops and netbooks, and prefer a traditional desktop metaphor. Although this is the 14.04 release, it can be considered a newer version than 14.10 as it has most of the updated packages.

**Ubuntu Server 14.10:** Ubuntu Server helps you make the most of your infrastructure. Along with other features, the latest release comes with the fastest, most secure bare-metal hypervisor and disk acceleration, extending SSD performance to large, cheap rotary disks.

**openSUSE 13.2 GNOME Live:** This is a free and Linux-based operating system for your PC, laptop or server. You can surf the Web, manage your e-mails and photos, do office work, play videos or music and have a lot of fun! This is a GNOME desktop that you can run from the DVD or from a USB stick and can also be installed as is, but it cannot be used to upgrade an existing system.

**openSUSE 13.2 KDE Live:** This is the Live KDE version of the latest release of openSUSE 13.2





# LED Lighting Is BIG Business

And it's going to get bigger!

**FREE**  
sample  
copies for  
B2B decision  
makers!

**Inviting the Who's Who  
in this business to stay  
updated on market trends  
by reading LED Bazaar**

LED Bazaar is a B2B publication which started as a 4-page section inside Electronics Bazaar (EB) magazine, and has now become a separate supplement along with EB. Shortly, it is expected to take shape of a separate magazine by itself.

To get a complimentary copy of LED Bazaar, simply register at <http://freeledb.efy.in>







**There is a place for Postgres in every datacenter.**

**WE CAN HELP YOU FIND IT.**

**Postgres delivers.** Enterprise class Performance, Security, Success.  
For 80% less.

*EnterpriseDB - The Postgres Database Company*

[Sales-india@enterprisedb.com](mailto:Sales-india@enterprisedb.com)



EnterpriseDB Software India Pvt Ltd  
Unit #3, Godrej Castlemaine  
Pune - 411 01 India  
T +91 20 30589500/01